

SHEAR-SHAPE COCYCLES FOR MEASURED LAMINATIONS AND ERGODIC THEORY OF THE EARTHQUAKE FLOW

AARON CALDERON AND JAMES FARRE

ABSTRACT. We extend Mirzakhani’s conjugacy between the earthquake and horocycle flows to a bijection, demonstrating conjugacies between these flows on all strata and exhibiting an abundance of new ergodic measures for the earthquake flow. The structure of our map indicates a natural extension of the earthquake flow to an action of the upper-triangular subgroup $P < \text{SL}_2\mathbb{R}$ and we classify the ergodic measures for this action as pullbacks of affine measures on the bundle of quadratic differentials. Our main tool is a generalization of the shear coordinates of Bonahon and Thurston to arbitrary measured laminations.

1. MAIN RESULTS

1.1. Conjugating earthquake and horocycle flow. This paper deals with two notions of unipotent flow over the moduli space \mathcal{M}_g of Riemann surfaces. The first is the *Teichmüller horocycle flow*, defined on the bundle $\mathcal{Q}^1\mathcal{M}_g$ of unit area quadratic differentials q by postcomposing the charts of the flat metric $|q|$ by the parabolic transformation $\begin{pmatrix} 1 & s \\ 0 & 1 \end{pmatrix}$. This flow is ergodic with respect to a finite measure induced by Lebesgue in local period coordinates [Mas82, Vee82] and is a fundamental object of study in Teichmüller dynamics.

The second is the *earthquake flow* on the bundle $\mathcal{P}^1\mathcal{M}_g$, whose fiber is the sphere of unit-length measured geodesic laminations on a hyperbolic surface. The earthquake flow is defined as a generalization of twisting about simple closed curves, or by postcomposing hyperbolic charts by certain piecewise-isometric transformations. While this flow is more mysterious, earthquakes are a familiar tool in Teichmüller theory, playing a central role in Kerckhoff’s proof of the Nielsen realization conjecture [Ker83], for example.

These two flows are both assembled from families of Hamiltonian flows (extremal length for horocycle [Pap86] and hyperbolic length for earthquake [Ker83, Wol83, SB01]) and exhibit similar non-divergence properties [MW02], but the horocycle flow belongs properly to the flat-geometric viewpoint and the earthquake flow to the hyperbolic one. All the same, in [Mir08, Theorem 1.1] Mirzakhani established a bridge between the two worlds, demonstrating a measurable conjugacy between the earthquake and horocycle flows. Consequently, the earthquake flow is ergodic with respect to the measure class of Lebesgue on $\mathcal{P}^1\mathcal{M}_g$.

In this article, we deepen this connection between flat and hyperbolic geometry, proving that the correspondence can be further upgraded to yield new results on both the ergodic theory of the earthquake flow and the structure of Teichmüller space.

Theorem A. Mirzakhani’s conjugacy extends to a bijection

$$\mathcal{O} : \mathcal{P}^1\mathcal{M}_g \leftrightarrow \mathcal{Q}^1\mathcal{M}_g$$

that conjugates earthquake flow to horocycle flow.

The moduli space of quadratic differentials is naturally partitioned into *strata* $\mathcal{Q}^1\mathcal{M}_g(\underline{\kappa})$, disjoint subsets parametrizing unit-area differentials with zeros of order $\underline{\kappa} = (\kappa_1, \dots, \kappa_n)$. Similarly, for any $\underline{\kappa}$ we may define the *regular locus* $\mathcal{P}^1\mathcal{M}_g^{\text{reg}}(\underline{\kappa})$ to be the set of (X, λ) where λ cuts X into ideal polygons with $(\kappa_1+2, \dots, \kappa_n+2)$ many sides, each with a cyclic symmetry of that order.

With this notation, Mirzakhani’s conjugacy can more precisely be stated as the existence of a bijection

$$\mathcal{P}^1\mathcal{M}_g^{\text{reg}}(1^{4g-4}) \leftrightarrow \mathcal{Q}^1\mathcal{M}_g^{\text{nsc}}(1^{4g-4})$$

taking earthquake flow to horocycle flow, where the superscript nsc specifies the (full-measure) sublocus of the stratum consisting of those differentials with no horizontal saddle connections.

One of our main applications of Theorem A is to produce an analogue of Mirzakhani’s conjugacy for components of strata (even those coming from global squares of Abelian differentials), confirming a conjecture of Alex Wright [Wri18, Remark 5.6] (see also [Wri20, Problems 12.5 and 12.6]).

Theorem B. For every $\underline{\kappa}$, the map \mathcal{O} restricts to a bijection

$$\mathcal{P}^1\mathcal{M}_g^{\text{reg}}(\underline{\kappa}) \leftrightarrow \mathcal{Q}^1\mathcal{M}_g^{\text{nsc}}(\underline{\kappa})$$

that takes earthquake to horocycle flow and (generalized) stretch rays to Teichmüller geodesics.

While strata of holomorphic quadratic differentials are generally not connected, for $g \neq 4$ their connected components are classified by whether or not they consist of squares of abelian differentials and the parity of the induced spin structure (both of which depend only on the horizontal foliation when there are no horizontal saddles), as well as hyperellipticity [KZ03, Lan08].¹ The bijection \mathcal{O} respects both the horizontal direction and the $\text{Mod}(S)$ action, so Theorem B can be refined to describe the preimages of these components.

As an immediate consequence of Theorem B, the earthquake flow is ergodic with respect to the pushforward by \mathcal{O}^{-1} of the Masur-Veech measure on any component of any stratum of quadratic differentials.

1.2. Geodesic flows and P -invariant measures. Pulling back the Teichmüller geodesic flow via \mathcal{O} allows us to specify a family of “dilation rays” which serve as a geodesic flow for the earthquake flow’s parabolic action and in many cases project to geodesics for Thurston’s Lipschitz asymmetric metric. Combining dilation rays and the earthquake flow therefore gives a action of the upper triangular subgroup $P < \text{SL}_2\mathbb{R}$ on $\mathcal{P}^1\mathcal{M}_g$ by “stretchquakes.” See Section 15.3.

Due in part to the failure of \mathcal{O} to be continuous, the stretchquake action on $\mathcal{P}^1\mathcal{M}_g$ is not by homeomorphisms but rather by measurable bijections. More precisely, it preserves the σ -algebra obtained by pulling back the Borel σ -algebra of $\mathcal{Q}^1\mathcal{M}(S)$ along \mathcal{O} . In a forthcoming sequel [CFA], the authors show that \mathcal{O} is actually a measurable isomorphism with respect to the Borel σ -algebra on $\mathcal{P}^1\mathcal{M}_g$ and that the stretchquake action restricted to each $\mathcal{P}^1\mathcal{M}_g^{\text{reg}}(\underline{\kappa})$ is by homeomorphisms; see also Remark 2.2.

Remark 1.1. In fact, Arana-Herrera and Wright have recently shown that there is *no* continuous map conjugating the earthquake flow to horocycle flow, at least when $\mathcal{P}^1\mathcal{M}_g$ and $\mathcal{Q}^1\mathcal{M}_g$ are equipped with their standard topologies [AHW22].

In their foundational work on $\text{SL}_2\mathbb{R}$ -invariant ergodic measures on the moduli space of flat surfaces, Eskin and Mirzakhani [EM18, Theorem 1.4] proved that the support of any P -invariant ergodic measure on $\mathcal{Q}^1\mathcal{M}_g$ is locally an affine manifold cut out by linear equations in period coordinates. Our conjugacy translates this classification into a classification of ergodic measures for the extension of the earthquake flow defined above:

Theorem C. Every stretchquake-invariant ergodic measure is the pullback of an affine measure.

Proof. If ν is a stretchquake-invariant ergodic measure on $\mathcal{P}^1\mathcal{M}_g$, then $\mathcal{O}_*\nu$ is a P -invariant ergodic measure on $\mathcal{Q}^1\mathcal{M}_g$, which is affine by [EM18, Theorem 1.4]. \square

Using this correspondence we obtain a geometric rigidity phenomenon for stretchquake-invariant ergodic measures on $\mathcal{P}^1\mathcal{M}_g$: the generic point is made out of a fixed collection of regular ideal polygons.

Corollary 1.2. For any stretchquake-invariant ergodic probability measure ν on $\mathcal{P}^1\mathcal{M}_g$, there is some $\underline{\kappa}$ so that ν -almost every (X, λ) lies in $\mathcal{P}^1\mathcal{M}_g^{\text{reg}}(\underline{\kappa})$.

This in particular implies that the dynamics of the stretchquake action with respect to any ergodic probability measure are measurably the same as its restriction to a stratum, on which we can identify dilation rays as (directed, unit-speed) geodesics for the Lipschitz asymmetric metric on $\mathcal{T}(S)$ (see Proposition 15.12).

Remark 1.3. We note that general ergodic measures for the stretchquake action can look quite different than the Lebesgue measure class on $\mathcal{P}^1\mathcal{M}_g$, even when pushed down to \mathcal{M}_g .

For example, if ν gives full measure to $\mathcal{P}^1\mathcal{M}_g^{\text{reg}}(4g-4)$ then a ν -generic point is obtained by gluing together a single regular ideal $(4g-2)$ -gon; in particular, the injectivity radius at the center of the polygon can be arbitrarily large, allowing $g \rightarrow \infty$. This implies that ν gives zero mass to (the restriction of $\mathcal{P}^1\mathcal{M}_g$ to) sufficiently thin parts of moduli space, as any (X, λ) where X has a very short pants decomposition has injectivity radius uniformly bounded above.

Remark 1.4. While an important result of [EM18] is that any P -invariant ergodic measure on $\mathcal{Q}^1\mathcal{M}_g$ is actually $\text{SL}_2\mathbb{R}$ -invariant, the circle action on $\mathcal{Q}^1\mathcal{M}_g$ (corresponding to rotating a quadratic differential) does not have an obvious geometric interpretation on $\mathcal{P}^1\mathcal{M}_g$. See also [Wri20, Problems 12.3 and 12.4].

¹In genus 4, there are certain strata whose components have only been characterized via algebraic geometry [CM14].

1.3. Dual foliations from hyperbolic structures. A foundational result of Gardiner and Masur (Theorem 2.1 below) states that quadratic differentials are parametrized by their real and imaginary parts; equivalently, their vertical and horizontal foliations (or laminations). In particular, the real-analytic submanifold $\mathcal{F}^{uu}(\lambda)$ of all quadratic differentials with horizontal lamination λ can be identified with the space $\mathcal{MF}(\lambda)$ of foliations that bind together with λ . See Section 2 for a formal definition. As the horocycle flow preserves the horizontal foliation, it induces a flow on $\mathcal{MF}(\lambda)$.

Mirzakhani's conjugacy and our extension therefore both follow from the construction of flow-equivariant maps that assign to a hyperbolic surface X and a measured lamination λ a “dual” measured foliation.

For maximal laminations λ , this dual is the *horocyclic foliation* $F_\lambda(X)$ introduced by Thurston [Thu98], obtained by foliating the spikes of each triangle of $X \setminus \lambda$ by horocycles and extending across the leaves of λ . The measure of an arc transverse to $F_\lambda(X)$ is then the total distance along λ between horocycles meeting the arc at endpoints. As $F_\lambda(X)$ necessarily binds S together with λ , this defines a map

$$F_\lambda : \mathcal{T}(S) \rightarrow \mathcal{MF}(\lambda).$$

We endow $\mathcal{MF}(\lambda)$ with the real-analytic structure coming from its identification with $\mathcal{F}^{uu}(\lambda)$. The main engine of Mirzakhani's conjugacy is the following theorem of Bonahon [Bon96] and Thurston [Thu98]; see also Section 2.1 for a discussion of her interpretation of this result.

Theorem 1.5 (Bonahon, Thurston). For any maximal λ , the horocyclic foliation map F_λ is a real-analytic homeomorphism which takes the earthquake in λ to the horocycle flow restricted to $\mathcal{MF}(\lambda) \cong \mathcal{F}^{uu}(\lambda)$ in a time-preserving way. Moreover, the family $\{F_\lambda\}$ is equivariant with respect to the $\text{Mod}(S)$ action. That is, $F_{g\lambda}(gX) = gF_\lambda(X)$ for all $g \in \text{Mod}(S)$.

When λ is not maximal the horocyclic foliation is no longer defined. The first thing one might try is to simply choose a completion of λ , but this approach is too naïve. Indeed, this would require choosing a completion of every lamination, which necessarily destroys $\text{Mod}(S)$ -equivariance because laminations (and differentials) can have symmetries.² Such a map will not descend to moduli space and is therefore unsuitable for our applications. Besides, for our purposes it is important that the geometry of the subsurfaces of $X \setminus \lambda$ predicts the singularity structure of the corresponding differential.

If one restricts their attention to the case when λ is filling and cuts X into regular ideal polygons then there is a canonical notion of horocyclic foliation. While this construction is equivalent on the regular locus to the more general procedure we describe just below, any attempt to prove Theorem B with this restricted viewpoint would necessarily rely on ($\text{Mod}(S)$ -equivariant) descriptions of the loci of surfaces built from regular polygons, as well as the intersection of $\mathcal{F}^{uu}(\lambda)$ with strata, results which (to the knowledge of the authors) were heretofore unknown. Compare Corollary 2.6 and Section 2.2.

We therefore place no restrictions on the topological type or the complementary geometry of λ . Following a suggestion of Yi Huang (communicated to us by Alex Wright), we prove that the correct analogue of the horocyclic foliation for non-maximal λ is the *orthogeodesic foliation* $\mathcal{O}_\lambda(X)$, whose leaves are the fibers of the closest point projection to λ and whose measure is given by length of the projection to λ . As in the maximal case, the orthogeodesic foliation binds together with λ , inducing a map

$$\mathcal{O}_\lambda : \mathcal{T}(S) \rightarrow \mathcal{MF}(\lambda).$$

See Section 5 for a more detailed discussion of this construction.

Theorem D. For any $\lambda \in \mathcal{ML}(S)$, the orthogeodesic foliation map \mathcal{O}_λ is a homeomorphism which takes the earthquake in λ to the horocycle flow restricted to $\mathcal{MF}(\lambda) \cong \mathcal{F}^{uu}(\lambda)$ in a time-preserving way. Moreover, the family $\{\mathcal{O}_\lambda\}$ is equivariant with respect to the $\text{Mod}(S)$ action. That is, $\mathcal{O}_{g\lambda}(gX) = g\mathcal{O}_\lambda(X)$ for all $g \in \text{Mod}(S)$.

Although $\mathcal{MF}(\lambda)$ does not have an obvious smooth structure, the map \mathcal{O}_λ still exhibits a surprising amount of regularity; see Theorem E.

The proof of Theorem D requires generalizing Bonahon's machinery of transverse cocycles to new combinatorial objects called “shear-shape cocycles” which capture the essential structure of the orthogeodesic foliation; see Section 2.1 just below. The space of shear-shape cocycles forms a common coordinatization of

²For example, take γ to be a simple closed curve; completions of γ correspond to triangulations of $X \setminus \gamma$ where the boundaries are shrunk to cusps (up to a choice of spiraling about each side of γ). The space of such triangulations carries a rich $\text{Stab}(\gamma)$ action, and a computation shows that the horocyclic foliations for two completions in the same $\text{Stab}(\gamma)$ orbit need not be equal.

both $\mathcal{T}(S)$ and $\mathcal{MF}(\lambda)$ that is compatible with the map \mathcal{O}_λ and reveals an abundance of structure encoded in the orthogeodesic foliation:

- When λ cuts X into regular ideal polygons, the orthogeodesic and horocyclic foliations agree.
- The locus of points of X which are closest to at least two leaves of λ forms a piecewise geodesic spine for $X \setminus \lambda$ which captures the geometry and topology of the complementary subsurfaces (see Theorem 6.4). Moreover, this spine is exactly the diagram of horizontal separatrices for the quadratic differential with horizontal foliation λ and vertical foliation $\mathcal{O}_\lambda(X)$.
- For every measure μ on λ , the intersection of μ and $\mathcal{O}_\lambda(X)$ is the hyperbolic length of μ on X .
- The pullback of Teichmüller geodesics with no horizontal saddle connections are geodesics with respect to Thurston's Lipschitz (asymmetric) metric (Proposition 15.12).

A statement similar to Theorem D is probably also true for any (unmeasured) geodesic lamination, but for technical reasons regarding compatibility of complementary subsurfaces and the spiraling behavior of λ we have restricted ourselves to the measured setting. See Remark 7.10.

The orthogeodesic foliation map can also be thought of as relating the hyperbolic and extremal length functions $\ell_\lambda(\cdot)$ and $\text{Ext}_\lambda(\cdot)$ for any fixed λ . Indeed, a seminal theorem of Hubbard and Masur [HM79] states that the natural projection

$$\pi : \mathcal{F}^{uu}(\lambda) \rightarrow \mathcal{T}(S)$$

that records only the complex structure underlying a differential is a homeomorphism. Combining this with the fact that the extremal length of λ on Y is exactly the area of the differential $\pi^{-1}(Y)$, we deduce that

Corollary 1.6. For every $\lambda \in \mathcal{ML}(S)$, the map $\pi \circ \mathcal{O}_\lambda$ is a $\text{Stab}(\lambda)$ -equivariant self-homeomorphism of $\mathcal{T}(S)$ that takes the hyperbolic length function $\ell_\lambda(\cdot)$ to the extremal length function $\text{Ext}_\lambda(\cdot)$.

1.4. Acknowledgments. The authors would firstly like to thank Alex Wright for providing the germ of this project and useful suggestions, as well as for helpful comments on a preliminary draft of this paper. We are grateful to Francis Bonahon, Feng Luo, Howie Masur, and Jing Tao for lending their expertise and for enlightening discussions.

The authors would also like to thank those who contributed helpful comments or with whom we had clarifying conversations, including Daniele Alessandrini, Francisco Arana-Herrera, Mladen Bestvina, Jon Chaika, Vincent Delecroix, Valentina Disarlo, Spencer Dowdall, Ben Dozier, Aaron Fenyes, Ser-Wei Fu, Curt McMullen, Mareike Pfeil, Beatrice Pozzetti, John Smillie, and Sam Taylor. We also want to thank the anonymous referee for their careful reading and useful comments.

Finally, the authors are indebted to Yair Minsky for his dedicated guidance and mentorship throughout all stages of this project, as well as his generous listening and insightful comments.

The first author gratefully acknowledges support from NSF grants DGE-1122492, DMS-161087, and DMS-2005328, and travel support from NSF grants DMS-1107452, -1107263, and -1107367 “RNMS: Geometric Structures and Representation Varieties” (the GEAR Network). The second author gratefully acknowledges support from NSF grants DMS-1246989, DMS-1509171, DMS-1902896, DMS-161087, and DMS-2005328.

Portions of this work were accomplished while the authors were visiting MSRI for the Fall 2019 program “Holomorphic Differentials in Mathematics and Physics,” and the authors would like to thank the venue for its hospitality and excellent working environment. Part of this material is based upon work supported by the National Science Foundation under Grant No. DMS-1928930 while the second author participated in a program hosted by the Mathematical Sciences Research Institute in Berkeley, California, during the Fall 2020 semester on “Random and Arithmetic Structures in Topology.” The second author would also like to thank the both the Department of Mathematics at the University of Utah and the Mathematics Institute at Universität Heidelberg for their hospitality and rich working environments.

CONTENTS

1. Main results	1
2. About the proof	4
3. Outline of the paper	10
4. Crowned hyperbolic surfaces	12
5. The orthogeodesic foliation	14
6. Cellulating crowned Teichmüller spaces	19

7. Transverse and shear-shape cocycles	24
8. The structure of shear-shape space	32
9. Train track coordinates for shear-shape space	36
10. Shear-shape coordinates for transverse foliations	40
11. Flat deformations in shear-shape coordinates	48
12. Shear-shape coordinates for hyperbolic metrics	50
13. Measuring hyperbolic shears and shapes	51
14. Shape-shifting cocycles	59
15. Shear-shape coordinates are a homeomorphism	77
16. Future and ongoing work	84
References	86
Index	89

2. ABOUT THE PROOF

Given Theorem D which associates to (X, λ) a dual foliation $\mathcal{O}_\lambda(X)$ describing the geometry of the pair, it is not difficult to prove Theorems A and B. First, we recall the relationship between differentials, foliations, and laminations in a little more detail.

The space of measured foliations (up to equivalence) on a closed surface S of genus $g \geq 2$ is denoted $\mathcal{MF}(S)$. There is a canonical identification [Lev83] between $\mathcal{MF}(S)$ and $\mathcal{ML}(S)$, the space of measured laminations on S ; throughout this paper we will implicitly pass between the two notions at will, depending on our situation. By \mathcal{QT}_g and $\mathcal{Q}^1\mathcal{T}_g$ we mean the bundle of holomorphic quadratic differentials over the Teichmüller space and the locus of unit area quadratic differentials, respectively. We similarly denote $\mathcal{PT}_g = \mathcal{T}(S) \times \mathcal{ML}(S)$ and $\mathcal{P}^1\mathcal{T}_g$ the locus of pairs (X, λ) where λ has unit length on X .

To every $q \in \mathcal{QT}_g$ one may associate the real measured foliation $|\operatorname{Re}(q)|$ which measures the total variation of the real part of the holonomy of an arc; the imaginary foliation $|\operatorname{Im}(q)|$ is defined similarly. These foliations have vertical, respectively horizontal, trajectories, and so we will also refer to them as the vertical and horizontal foliations (or laminations) of q and write

$$q = q(|\operatorname{Re}(q)|, |\operatorname{Im}(q)|).$$

A foundational theorem of Gardiner and Masur implies that the real and imaginary foliations completely determine q , and that given any two foliations which “fill up” the surface, one can integrate against their measures to recover a quadratic differential.

A pair of measured foliations/laminations (η, λ) is said to *bind* S if for every $\gamma \in \mathcal{ML}(S)$,

$$i(\gamma, \eta) + i(\gamma, \lambda) > 0,$$

where $i(\cdot, \cdot)$ is the geometric intersection pairing. In the literature, such pairs are sometimes called *filling*, though we choose to distinguish the topological notion of filling from the measure-theoretic notion of binding.

Theorem 2.1 ([GM91, Theorem 3.1]). There is a $\operatorname{Mod}(S)$ -equivariant homomorphism

$$\mathcal{QT}(S) \cong \mathcal{MF}(S) \times \mathcal{MF}(S) \setminus \Delta$$

where Δ is the set of all non-binding pairs (η, λ) . In particular, the set $\mathcal{F}^{uu}(\lambda)$ of all quadratic differentials with $|\operatorname{Im}(q)| = \lambda$ may be identified with $\mathcal{MF}(\lambda)$, the set of foliations which together bind with λ .

Proof of Theorems A and B. By definition, there is a $\operatorname{Mod}(S)$ -equivariant projection $\mathcal{PT}_g \rightarrow \mathcal{ML}(S)$ with fiber $\mathcal{T}(S)$. Theorem 2.1 implies there is a $\operatorname{Mod}(S)$ -equivariant projection $\mathcal{QT}_g \rightarrow \mathcal{ML}(S)$ whose fiber over λ may be identified with $\mathcal{MF}(\lambda)$. Applying Theorem D on the fibers therefore yields an equivariant bijection

$$\mathcal{O} : \mathcal{PT}_g \leftrightarrow \mathcal{QT}_g$$

which takes unit-length laminations to unit-area differentials (Corollary 13.14), and quotienting by the $\operatorname{Mod}(S)$ action proves Theorem A.

Furthermore, we observe that the spine of the orthogeodesic foliation of a regular ideal $(k+2)$ -gon is just a star with $k+2$ edges, which corresponds to the separatrix diagram of a zero of order k when there are no horizontal saddle connections. Thus \mathcal{O} restricts to the promised conjugacy on strata (Theorem B). \square

Remark 2.2. Mirzakhani's conjugacy is defined on the Borel subset $\mathcal{PT}_g^{\text{reg}}(1^{4g-4}) \subset \mathcal{PT}_g$ of full Lebesgue measure and is moreover Borel measurable on its domain of definition. The latter assertion is a consequence of a stronger result, namely that $\mathcal{PT}_g^{\text{reg}}(1^{4g-4}) \rightarrow \mathcal{QT}_g$ is continuous (with respect to the subspace topology on $\mathcal{PT}_g^{\text{reg}}(1^{4g-4})$).

While convergence of measured laminations (in measure) does not typically imply Hausdorff convergence of the supports, whenever a sequence $\{\lambda_n\}$ of maximal measured laminations converges to a maximal measured lamination λ , then λ_n is eventually carried (snugly) on a maximal train track also carrying λ . From here, it is not difficult to deduce that $\lambda_n \rightarrow \lambda$ in the Hausdorff topology [ZB04] and thus the horocyclic foliations $F_{\lambda_n}(X)$ converge to $F_{\lambda}(X)$. Intuitively, the leaves of λ_n intersect the leaves of λ with small angle (depending on the specific surface on which they are realized), so the orthogonal directions become more parallel.

In forthcoming work [CFa], the authors extend these ideas and prove that \mathcal{O} is (everywhere) Borel measurable with Borel measurable inverse by identifying a countable partition of \mathcal{PT}_g and \mathcal{QT}_g into Borel subsets on which \mathcal{O} is homeomorphic. See also Section 16.

In general, the compact edges of the spine of a pair (X, λ) correspond exactly to horizontal saddle connections in the differential $\mathcal{O}(X, \lambda)$. This observation allows us to prove that the generic point for a P -invariant ergodic probability measure on $\mathcal{P}^1\mathcal{M}_g$ consists of pairs (X, λ) where λ cuts X into a fixed set of regular ideal polygons.

Proof of Corollary 1.2. Using our conjugacy, the desired statement is equivalent to the fact that any P -invariant ergodic probability measure on $\mathcal{Q}^1\mathcal{M}_g$ is (a) supported in a single stratum and (b) gives 0 measure to the set of differentials with horizontal saddle connections.

The first statement is implied by ergodicity, while the second follows from the fact that the measure is actually $\text{SL}_2\mathbb{R}$ -invariant [EMM15]. Indeed, for any quadratic differential q , the Lebesgue measure of the set of directions θ such that $e^{i\theta}q$ has a saddle connection is 0, so Fubini's theorem implies (b). \square

Refining the proof by considering connected components of strata, we see that we can also conclude that ν -almost every pair has the same orientability, spin, and hyperellipticity properties.

2.1. Shear-shape coordinates. Our strategy to prove Theorem D follows Mirzakhani's interpretation of Theorem 1.5 in which she clarifies the relationship between Thurston's geometric perspective on the horocyclic foliation and Bonahon's powerful analytic approach in terms of transverse cocycles. Namely, she shows that the horocyclic foliation map F_λ is compatible with shearing coordinates for both hyperbolic structures and measured foliations. To motivate our construction, we give a brief outline of Mirzakhani's proof below.

A (real-valued) *transverse cocycle* for λ is a finitely additive signed measure on arcs transverse to λ that is invariant under isotopy transverse to λ ; observe that transverse measures are themselves transverse cocycles. These objects equivalently manifest as transverse Hölder distributions, cohomology classes, or weight systems on snug train tracks [Bon97a, Bon96, Bon97b]. The space $\mathcal{H}(\lambda)$ of transverse cocycles forms a finite dimensional vector space which carries a natural homological intersection pairing which is non-degenerate when λ is maximal. The intersection pairing then identifies a “positive locus” $\mathcal{H}^+(\lambda) \subset \mathcal{H}(\lambda)$ cut out by finitely many geometrically meaningful linear inequalities. See also Section 7.1.

In [Bon96, Theorem A], Bonahon proved that for any maximal geodesic lamination λ there is a real-analytic homeomorphism $\sigma_\lambda : \mathcal{T}(S) \rightarrow \mathcal{H}^+(\lambda)$ that takes a hyperbolic metric to its “shearing cocycle,” which essentially records the signed distance along λ between the centers of ideal triangles in the complement of λ . Mirzakhani then constructed a homeomorphism I_λ (essentially by a well-chosen system of period coordinates) that coordinatizes $\mathcal{MF}(\lambda)$ by $\mathcal{H}^+(\lambda)$ and for which the following diagram commutes [Mir08, §§5.2, 6.2]:

$$(1) \quad \begin{array}{ccc} \mathcal{T}(S) & \xrightarrow{F_\lambda} & \mathcal{MF}(\lambda) \\ \sigma_\lambda \searrow & & \swarrow I_\lambda \\ & \mathcal{H}^+(\lambda) & \end{array}$$

Since $F_\lambda = I_\lambda^{-1} \circ \sigma_\lambda$ is a composition of homeomorphisms, it is itself a homeomorphism. As the construction of the horocyclic foliation requires no choices, the family $\{F_\lambda\}$ is necessarily $\text{Mod}(S)$ -equivariant. Finally, a direct computation shows that σ_λ transports the earthquake in λ to translation in $\mathcal{H}^+(\lambda)$ by λ , and I_λ similarly takes horocycle to translation, demonstrating Theorem 1.5.

Shear-shape cocycles. When λ is not maximal, the space of transverse cocycles is no longer suitable to coordinatize hyperbolic structures (or transverse foliations). Indeed, in this case the vector space $\mathcal{H}(\lambda)$ has dimension less than $6g - 6$ and its intersection form may be degenerate; this is a consequence of the fact that the Teichmüller space of $S \setminus \lambda$ now has a rich analytic structure that transverse cocycles cannot see.

In order to imitate Diagram (1) and its concomitant arguments for arbitrary $\lambda \in \mathcal{ML}(S)$, we therefore introduce the notion of *shear-shape cocycles* on λ . Roughly, a shear-shape cocycle consists of finitely additive signed data on certain arcs transverse to λ together with a weighted arc system that cuts $S \setminus \lambda$ into cells; this pair is also required to satisfy a certain compatibility condition mimicking features of the orthogeodesic foliation (Definition 7.12). Generalizing results of Luo [Luo07, Theorem 1.2 and Corollary 1.4], we show that such an arc system is equivalent to a hyperbolic structure on $S \setminus \lambda$ (Theorem 6.4), so shear-shape cocycles may equivalently be thought of as transverse data together with a compatible hyperbolic structure on the complementary subsurface(s). Like transverse cocycles, shear-shape cocycles also admit realizations as cohomology classes or weight systems on certain train tracks (Definition 7.5 and Proposition 9.5).

Remark 2.3. We note that only certain classes of arcs admit consistent weights when measured by a shear-shape cocycle, whereas transverse cocycles provide a measure to any arc transverse to λ . While this subtlety is exactly what allows us to understand how to relate shear-shape cocycles with the geometry of complementary subsurfaces, it also presents a number of technical challenges throughout the paper.

Unlike transverse cocycles, the space $\mathcal{SH}(\lambda)$ of shear-shape cocycles is not a vector space, instead forming a principal $\mathcal{H}(\lambda)$ bundle over a contractible analytic subvariety of $\mathcal{T}(S \setminus \lambda)$ (Theorem 8.1). All the same, the cohomological realization of shear-shape cocycles equips $\mathcal{SH}(\lambda)$ with an intersection form

$$\omega_{\mathcal{SH}} : \mathcal{SH}(\lambda) \times \mathcal{H}(\lambda) \rightarrow \mathbb{R}$$

that identifies a “positive locus” $\mathcal{SH}^+(\lambda)$ and equips both $\mathcal{SH}(\lambda)$ and $\mathcal{SH}^+(\lambda)$ with piecewise-integral-linear structures. The positive locus forms a $\mathcal{H}^+(\lambda)$ cone-bundle over the same subvariety of $\mathcal{T}(S \setminus \lambda)$ (Proposition 8.5) and fits into a familiar-looking commutative diagram:

$$(2) \quad \begin{array}{ccc} \mathcal{T}(S) & \xrightarrow{\sigma_\lambda} & \mathcal{MF}(\lambda) \\ & \searrow \sigma_\lambda & \swarrow I_\lambda \\ & \mathcal{SH}^+(\lambda) & \end{array}$$

where σ_λ and I_λ record shearing data along λ as well as shape data in the complementary subsurfaces. These maps can be thought of as a common generalization of Bonahon and Mirzakhani’s shear coordinates as well as Fenchel–Nielsen and Dehn–Thurston coordinates adapted to a pants decomposition (see Section 2.2). In the case when λ is orientable, the map I_λ can also be viewed as an extension of Minsky and Weiss’s description of the set of Abelian differentials with given horizontal foliation [MW14, Theorem 1.2].³

The conjugacy of Theorem D is then a consequence of the following structural theorem, which is an amalgam of the main technical results of the paper (compare Theorems 10.15, 12.1, and 13.13).

Theorem E. For any measured lamination λ , Diagram (2) commutes and all arrows are $\text{Stab}(\lambda)$ –equivariant homeomorphisms. Moreover,

- σ_λ is (stratified) real-analytic and transports the earthquake flow to translation by λ and the hyperbolic length of λ to $\omega_{\mathcal{SH}}(\cdot, \lambda)$.
- The weighted arc system underlying $\sigma_\lambda(X)$ records the hyperbolic structure $X \setminus \lambda$ under the correspondence of Theorem 6.4.
- I_λ is piecewise-integral-linear and transports horocycle flow to translation by λ and intersection with λ to $\omega_{\mathcal{SH}}(\cdot, \lambda)$.
- The weighted arc system underlying $I_\lambda(\eta)$ records the compact horizontal separatrices of $q(\eta, \lambda)$.

In the course of our proof, we also describe new “shape-shifting deformations” of hyperbolic surfaces which generalize Bonahon and Thurston’s cataclysms by shearing along a lamination while also varying the hyperbolic structures on complementary pieces. See Section 15.1.

³Technically, [MW14] investigates the family of Abelian differentials with a fixed horizontal foliation and fixed topological type of horizontal separatrix diagram, whereas our map applies to quadratic differentials (whether or not they are globally the square of an Abelian differential) and packages together all possible types of separatrix diagrams.

One particularly interesting family of deformations is obtained by dilation. The space $\mathcal{SH}^+(\lambda)$ admits a natural scaling action by $\mathbb{R}_{>0}$, and since both earthquake and horocycle flow are carried to translation in coordinates, this scaling action indicates extensions of each to P actions. A quick computation (Lemma 11.1) shows that the pullback of a dilation ray by I_λ is (a variant of) the Teichmüller geodesic flow, so the P action on the flat side is just the standard P action on \mathcal{QT}_g .

On the hyperbolic side, these dilation rays define our extension of the earthquake flow, and correspond to families of hyperbolic metrics on which the length of λ is scaled by a uniform factor. They are therefore natural candidates for (directed, unit-speed) geodesics for the Lipschitz asymmetric metric on $\mathcal{T}(S)$, and in some cases we can identify them as such (see Propositions 15.12 and 15.18, as well as Remarks 15.19 and 15.14).

Remark 2.4. Over the course of the paper we formalize the notion that shear-shape coordinates for hyperbolic structures are essentially the “real part” of period coordinates for \mathcal{PT}_g . Interpreting $\sigma_\lambda(X) + i\lambda$ as a complex weight system on a train track, Theorem C implies that the support of every stretchquake-invariant ergodic measure on $\mathcal{P}^1\mathcal{M}_g$ is locally an affine measure in train track charts. See Lemma 10.10.

Coordinatizing horospheres. Since the Thurston intersection form $\omega_{\mathcal{SH}}$ captures both the hyperbolic length of and geometric intersection with λ , the coordinate systems of Theorem E also allow us to give global descriptions of the level sets of these functions. In particular, we can recover Gardiner and Masur’s description of extremal length horospheres [GM91, p.236] as well as Bonahon’s description of the hyperbolic length ones (which is implicit in the structure of shear coordinates for maximal completions).

Corollary 2.5. Suppose that λ supports k ergodic transverse measures $\lambda_1, \dots, \lambda_k$. Then for all $L_1, \dots, L_k \in \mathbb{R}_{>0}$, the level sets

$$\{X \in \mathcal{T}(S) \mid \ell_X(\lambda_i) = L_i \text{ for all } i\} \text{ and } \{\eta \in \mathcal{MF}(\lambda) \mid i(\eta, \lambda_i) = L_i \text{ for all } i\}$$

are both homeomorphic to \mathbb{R}^{6g-6-k} .

Analyzing this coordinatization more closely, we see that in fact both level sets can be described as affine bundles of dimension $\dim_{\mathbb{R}} \mathcal{H}^+(\lambda) - k$ over the same subvariety of $\mathcal{T}(S \setminus \lambda)$ as underlies $\mathcal{SH}(\lambda)$.

From this refinement we are able to describe the intersection of the leaf $\mathcal{F}^{uu}(\lambda)$ with strata. The decomposition of period coordinates into real and imaginary parts shows that this intersection (when not empty) is locally homeomorphic to \mathbb{R}^d , where d is the complex dimension of the stratum; our work shows that these local homeomorphisms patch together to a global one. Compare [MW14, Theorem 1.2].

Corollary 2.6. Suppose that λ is a filling measured lamination that cuts a surface into polygons with $\kappa_1 + 2, \dots, \kappa_n + 2$ many sides, and let $\varepsilon = +1$ if λ is orientable and -1 otherwise. Let $\mathcal{QT}_g(\underline{\kappa}; \varepsilon)$ denote the union of the components of the stratum $\mathcal{QT}_g(\underline{\kappa}) \subset \mathcal{QT}_g$ that either are ($\varepsilon = +1$) or are not ($\varepsilon = -1$) global squares of Abelian differentials. Then

$$\{q \in \mathcal{QT}_g(\underline{\kappa}; \varepsilon) \mid |\text{Im}(q)| = \lambda\} \cong \mathcal{H}^+(\lambda) \cong \mathbb{R}^d$$

where d is the complex dimension of $\mathcal{QT}_g(\underline{\kappa}; \varepsilon)$.

Proof. Theorem E indicates that the metric graph of compact horizontal separatrices of $q(\eta, \lambda)$ is encoded by the weighted arc system underlying $I_\lambda(\eta)$. These weighted arc systems are organized in a piecewise-linear subvariety $\mathcal{B}(S \setminus \lambda)$ of a product of *weighted filling arc complexes* that encode the combinatorics of how a zero of order κ_i can split up into lower order zeros joined by horizontal saddle connections (see Sections 6, 7.3, 10.1, and Figure 5). For differentials in the indicated set, there are no compact horizontal separatrices, and so the underlying arc system is always the *empty* (filling) arc system $\emptyset \in \mathcal{B}(S \setminus \lambda)$. In other words, the image of $\{q \in \mathcal{QT}_g(\underline{\kappa}; \varepsilon) \mid |\text{Im}(q)| = \lambda\}$ in coordinates is just the fiber over \emptyset , where Proposition 8.5 identifies $\mathcal{SH}^+(\lambda)$ as an $\mathcal{H}^+(\lambda)$ -bundle over $\mathcal{B}(S \setminus \lambda)$.

The second isomorphism $\mathcal{H}^+(\lambda) \cong \mathbb{R}^d$ is just a dimension count (see Lemmas 4.6 and 7.3 in particular). \square

In general, we see that $\mathcal{F}^{uu}(\lambda) \cap \mathcal{QT}_g(\underline{\kappa}; \varepsilon)$ forms a $\mathcal{H}^+(\lambda)$ bundle over a union of faces of an arc complex of $S \setminus \lambda$. As a consequence we find that the only obstruction to completeness of any such leaf comes from zeros colliding along a horizontal saddle connection (see also [MW14, Theorem 11.2]). This global description of $\mathcal{F}^{uu}(\lambda) \cap \mathcal{QT}_g(\underline{\kappa}; \varepsilon)$ also allows the import of arguments from homogeneous dynamics to investigate equidistribution in both $\mathcal{Q}^1\mathcal{M}_g$ and $\mathcal{P}^1\mathcal{M}_g$ and their strata; see the discussion in Section 16.

2.2. Generalized Fenchel–Nielsen coordinates. Our shear-shape coordinates for hyperbolic structures can be thought of as interpolating between the classical Fenchel–Nielsen coordinates adapted to a pants decomposition and Bonahon and Thurston’s shear coordinates. In both cases, one remembers the shapes of the complementary subsurfaces (pairs of pants and ideal triangles, respectively) and the space of all hyperbolic structures with given complementary shapes is parametrized by gluing data (twist/shear parameters).

For general λ , there is a map

$$\text{cut}_\lambda : \mathcal{T}(S) \rightarrow \mathcal{T}(S \setminus \lambda)$$

that remembers the induced hyperbolic structure on each complementary subsurface. Theorem 12.1 then implies that the image of cut_λ is a real-analytic subvariety $\mathcal{B}(S \setminus \lambda)$ of $\mathcal{T}(S \setminus \lambda)$ consisting of those structures satisfying a “metric residue condition” (see Lemma 13.1). In the case where each component of λ is either non-orientable or a simple closed curve, $\mathcal{B}(S \setminus \lambda)$ is just the space of hyperbolic structures for which the two boundary components of the cut surface corresponding to a simple curve component of λ have equal length. Theorem 12.1 together with the structure of $\mathcal{SH}^+(\lambda)$ also allows us to identify the fiber $\text{cut}_\lambda^{-1}(Y)$ over any $Y \in \mathcal{B}(S \setminus \lambda)$ with the gluing data $\mathcal{H}^+(\lambda)$ ⁴ (though not in a canonical way).

We summarize this discussion in the following triptych:

$$(3) \quad \begin{array}{ccc} \text{Fenchel–Nielsen} & \text{Shear-shape} & \text{Shear} \\ \mathbb{R}^{3g-3} \longrightarrow \mathcal{T}(S) & \mathcal{H}^+(\lambda) \longrightarrow \mathcal{T}(S) & \mathcal{H}^+(\lambda) \longrightarrow \mathcal{T}(S) \\ \downarrow & \downarrow & \downarrow \\ \mathbb{R}_{>0}^{3g-3} & \mathcal{B}(S \setminus \lambda) & \{\text{pt}\} \\ \lambda \text{ a pants decomposition} & \lambda \text{ arbitrary} & \lambda \text{ maximal} \end{array}$$

In each coordinate system, $\mathcal{T}(S)$ is the total space of a fiber bundle over a base space of allowable shape data on the subsurface complementary to λ , while the fiber consists of gluing data.

A completely analogous picture also holds for foliations transverse to λ , demonstrating I_λ as a common generalization of both Dehn–Thurston and Mirzakhani’s shear coordinates.

2.3. Fenchel–Nielsen and Dehn–Thurston via shears and shapes. In order to give the reader a concrete example of shear-shape coordinates, we include here a discussion of our construction for $\lambda = P$ a pants decomposition. In this case, we see that shear-shape coordinates are just a (mild) reformulation of the classical Fenchel–Nielsen and Dehn–Thurston ones.

First we consider a hyperbolic structure X . A pair of pants in $X \setminus P$ is typically parametrized by its boundary lengths (a, b, c) , or equivalently, by the alternating side lengths of either of the right angled hexagons coming from cutting along seams. The orthogeodesic foliation on a pair of pants picks out either a pair or triple of seams (those which are realized as leaves of $\mathcal{O}_P(X)$), each weighted by the length of a boundary arc consisting of endpoints of leaves of $\mathcal{O}_P(X)$ isotopic to the seam. See Figure 1. In this case, these lengths are just simple (piecewise) linear combinations of the boundary lengths and the metric residue condition defining $\mathcal{B}(S \setminus P)$ just states that the boundaries that are glued together must have the same length. See Figure 1.

The space $\mathcal{H}^+(P)$ reduces to a sum of the twist spaces for each curve of P , and so Theorem 8.5 implies that $\mathcal{SH}^+(P)$ is a principal \mathbb{R}^{3g-3} bundle over $\mathcal{B}(S \setminus P) \cong \mathbb{R}_{>0}^{3g-3}$. The transverse data recorded by this twist space then describes the signed distance between certain reference points in pairs of right-angled hexagons in \tilde{X} that are adjacent to the same curve of \tilde{P} , which is the same as the twist parameter measured by the appropriate choice of Fenchel–Nielsen coordinates.⁵

We can similarly recognize $I_\lambda : \mathcal{MF}(P) \rightarrow \mathcal{SH}^+(P)$ as Dehn–Thurston coordinates. Note first that there are no essential simple closed curves in the complement of P , so $\mathcal{MF}(P)$, the space of measured foliations that intersect every curve in the support of P , is the same as the space of measured foliations not contained in the support of P . Now from any integral point $\sigma \in \mathcal{SH}^+(P)$ we can construct a multicurve α with prescribed intersection and twisting parameters as follows: the weighted arc system describes how strands of α pass between and meet the components of P , while the transverse data recorded by $\mathcal{H}^+(P) \cong \mathbb{R}^P$

⁴See the discussion around 18 in regards to the positivity condition for disconnected λ ; in essence, $\mathcal{H}^+(\lambda)$ is the product of $\mathcal{H}^+(\lambda_i)$ for each non-closed minimal component together with the twisting data around simple closed curves.

⁵Fenchel–Nielsen coordinates always involve some choice of section of the space of twists over the length parameters, and so have only the structure of a principal \mathbb{R}^{3g-3} bundle over \mathbb{R}_+^{3g-3} .

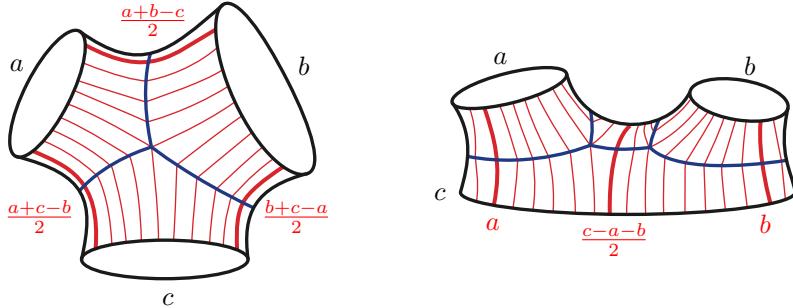


FIGURE 1. The orthogeodesic foliation on pairs of pants. Note that the weight of each bolded arc is a linear combination of the boundary lengths, hence the correspondence between shear-shape and Fenchel–Nielsen/Dehn–Thurston coordinates. If any of the weights is zero, the orthogeodesic foliation only picks out the two seams with non-zero weights.

describes the extent that strands of α wrap around components of P . This procedure is clearly reversible and can easily be extended to transverse foliations using a family of standard train tracks on each pair of pants (see [PH92, §2.6]). As in the hyperbolic case, one can easily pass between these coordinates and the standard Dehn–Thurston ones just by replacing the count of strands of α going from one boundary to the other with the total intersection of α with each boundary.

3. OUTLINE OF THE PAPER

The rest of this paper is roughly divided into four parts, corresponding to the orthogeodesic foliation, shear-shape cocycles, and shear-shape coordinates for flat and hyperbolic structures, as well as a collection of further directions for investigation (Section 16). While the constructions of I_λ and σ_λ both rely on foundational results established in the first two parts, we have attempted to direct the reader eager to understand our coordinates to the most important statements of these sections.

We expect that the reader is familiar with many of the standard constructions of Teichmüller theory, as well as the definitions of both the earthquake and horocycle flows; we recommend [MW02, Section 4] for a particularly lucid overview of the relevant objects. We also refer the reader to [CB88] and [Thu82, Section 8] for more on laminations and to [PH92] for a comprehensive introduction to train tracks.

§§4–6: The orthogeodesic foliation. Cutting along a lamination results in a (possibly disconnected) hyperbolic surface Σ with crown boundary, and in Section 4 we recall some useful information about the Teichmüller spaces of such surfaces. One particularly important definition is that of the “metric residue” of a crown end, which is a generalization of boundary length and plays an important role in cohomological constraints on the shape data of shear-shape cocycles (Lemma 7.9).

With these preliminaries established, in Section 5 we discuss in more detail the orthogeodesic foliation and the hyperbolic geometry of X in a neighborhood of λ . In this section we also give a geometric interpretation of the map in Corollary 1.6 that relates hyperbolic and extremal length.

The most important result of this part occupies Section 6, in which we show that the orthogeodesic foliation restricted to Σ completely determines its hyperbolic structure. More explicitly, dual to each compact edge of the spine of $\mathcal{O}_\lambda(X)$ is a packet of properly isotopic arcs joining non-asymptotic boundary components of Σ . By assigning geometric weights to each of these packets we can therefore combinatorialize the restriction of $\mathcal{O}_\lambda(X)$ to Σ by a weighted, filling arc system.

Using a geometric limit argument, in Theorem 6.4 we prove that the map which associates to a hyperbolic structure on Σ the associated arc system is a $\text{Mod}(\Sigma)$ -equivariant stratified real-analytic homeomorphism between $\mathcal{T}(\Sigma)$ and a certain type of arc complex for Σ , generalizing a theorem of Luo [Luo07] for surfaces with totally geodesic boundary (see also [Mon09b, Do08, Ush99]). Moreover, by construction this map records both the combinatorial structure of the spine of $\mathcal{O}_\lambda(X)$ as well as the metric residue of the crowns of Σ .

Theorem 6.4 is used extensively throughout the paper in order to pass between the combinatorial data of a weighted arc system, the restriction of $\mathcal{O}_\lambda(X)$ to Σ , and the corresponding hyperbolic structure on Σ .

The proof is independent of the main line of argument; as such, the reader is encouraged to understand the statement, but may wish only to skim the proof.

§§7–9: The space of shear-shape cocycles. The second part of the paper is devoted to our construction of shear-shape cocycles for a given λ and an analysis of the space $\mathcal{SH}(\lambda)$ of all shear-shape cocycles. Upon reaching this section, the reader may find it useful to glance ahead to either Section 10 or 13 to instantiate our definitions.

After reviewing structural results on transverse cocycles, in Section 7 we give both cohomological and axiomatic definitions of shear-shape cocycles (Definitions 7.5 and 7.12, respectively), both predicated on some underlying weighted arc system on Σ . In Proposition 7.14 we prove these definitions agree. Using the cohomological description, we observe a constraint on the weighted arc systems that can underlie a shear-shape cocycle coming from metric residue conditions (Lemma 7.9); this can also be thought of as a generalization of the fact that one can only glue together totally geodesic boundary components of the same length (compare Lemma 13.1).

Letting $\mathcal{B}(S \setminus \lambda)$ denote the subvariety of the filling arc complex of Σ cut out by the aforementioned residue conditions, we show in Section 8 that the space $\mathcal{SH}(\lambda)$ of shear-shape cocycles forms a bundle of transverse cocycles over $\mathcal{B}(S \setminus \lambda)$ with some additional structure (Theorem 8.1) whose total space is a cell of dimension $6g - 6$ (Corollary 8.2). In this section we also introduce the Thurston intersection form on $\mathcal{SH}(\lambda)$ (Section 8.2) and prove that the positive locus $\mathcal{SH}^+(\lambda)$ it defines is itself a bundle over $\mathcal{B}(S \setminus \lambda)$ (Proposition 8.5).

Finally, in Section 9 we give train track coordinates for the space of shear-shape cocycles. The train tracks we use give a preferred decomposition of arcs on S into pieces that are measurable by shear-shape cocycles and as such give a useful way of specifying shear-shape cocycles by a finite amount of data. The weight space for a train track is also a natural model in which to consider local deformations of a shear-shape cocycle, a feature which we exploit in Section 14. In Section 9.3 we discuss how the piecewise integral linear structure induced by train track charts endows $\mathcal{SH}^+(\lambda)$ with a well defined integer lattice and preferred measure in the class of Lebesgue.

The reader willing to accept the structure theorems can adequately navigate the remaining two parts of the paper using weight systems on (augmented) train tracks as a local description of the structure of shear-shape space.

§§10 and 11: Coordinates for transverse foliations. At this point, we have established the structure necessary to coordinatize foliations transverse to λ by shear-shape cocycles.

A measured foliation $\eta \in \mathcal{MF}(\lambda)$ determines a holomorphic quadratic differential $q = q(\eta, \lambda) \in \mathcal{F}^{uu}(\lambda)$ via Theorem 2.1, and we begin by specifying an arc system $\underline{\alpha}(q)$ that records the horizontal separatrices of q . We then build a train track τ carrying λ from a triangulation by saddle connections (Construction 10.4); augmenting τ by the arc system $\underline{\alpha}(q)$ then allows us to realize the periods of the triangulation as a (cohomological) shear-shape cocycle $I_\lambda(\eta)$. This identification also gives a useful formula for $I_\lambda(\eta)$ as a weight system on the augmented train track τ (Lemma 10.10).

We then show that one can rebuild q just from the train track weights defined by $I_\lambda(\eta)$; a similar (but more technical) argument then gives that $I_\lambda(\eta) \in \mathcal{SH}^+(\lambda)$ (Proposition 10.12). This reconstruction technique together with the structure of shear-shape space therefore allows to deduce that I_λ is a homeomorphism onto its image. At the end of this section, we explain how the work done in the fourth and final part of the paper implies that I_λ surjects onto $\mathcal{SH}^+(\lambda)$ (Theorem 10.15), and why we choose to prove surjectivity this way. See Remark 10.16 in particular.

Since I_λ essentially yields period coordinates, it is not surprising that (a variant of) Teichmüller geodesic flow is given in coordinates by dilation (Lemma 11.1), while the Teichmüller horocycle flow is translation by λ (Lemma 11.2). We also naturally recover the “tremor deformations” introduced in [CSW20] as translation by measures μ supported on λ that are not necessarily absolutely continuous with respect to λ (Definition 11.3). Figure 17 details a dictionary between the language of [CSW20] and our own.

§§12–15: Coordinates for hyperbolic structures. In the final part of the paper, we use the geometry of the orthogeodesic foliation to coordinatize hyperbolic structures via shear-shape cocycles.

From Theorem 6.4, we know that the combinatorialization of $\mathcal{O}_\lambda(X)$ on each subsurface $S \setminus \lambda$ by a weighted arc system completely encodes the geometry of the pieces. Cutting $X \setminus \lambda$ further along the orthogeodesic realization of each such arc, we obtain a family of (partially ideal) right-angled polygons. The orthogeodesic

foliation equips each polygon with a natural family of basepoints, one on each of its sides adjacent to λ , that vary analytically in $\mathcal{T}(S \setminus \lambda)$. We are thus able to define a “shear” parameter between (some pairs of) degenerate polygons, and this shear data assembles together with the “shape” data on each subsurface to give instructions for gluing the polygonal pieces back together to obtain X .

In Section 12 we state the main Theorem 12.1, that the shear-shape coordinate map $\sigma_\lambda : \mathcal{T}(S) \rightarrow \mathcal{SH}^+(\lambda)$ is a homeomorphism, supply an outline of its proof, and derive some immediate corollaries. The construction of σ_λ is given in Section 13, where we formalize the discussion from the previous paragraph. We also prove that the central Diagram (2) commutes (Theorem 13.13), which then implies that σ_λ takes hyperbolic length to the Thurston intersection form (Corollary 13.14).

Section 14 is the most technical part of the paper. In it, we define the “shape-shifting” cocycles (Proposition 14.26) along which a hyperbolic structure can be deformed (Theorem 15.1); these deformations are generalizations of Thurston’s cataclysms or Bonahon’s shear deformations. Although the construction of a shape-shifting deformation is rather involved, we attempt to keep the reader informed of the geometric intuition that guides the construction throughout. Finally, in Section 15 we assemble all of the necessary ingredients to prove Theorem 12.1. That the earthquake along λ is given by translation by λ in $\mathcal{SH}^+(\lambda)$ (Corollary 15.2) is an immediate consequence of the construction of shape-shifting deformations as generalizations of cataclysms. We then discuss how the action of dilation in coordinates can sometimes be identified with directed geodesics in Thurston’s asymmetric metric (Propositions 15.12 and 15.18).

4. CROWNED HYPERBOLIC SURFACES

When a hyperbolic surface is cut along a geodesic multicurve, the (completion of the) resulting space is a compact hyperbolic surface with compact, totally geodesic boundary. When the same surface is cut along a geodesic lamination, the (completion of the) complementary subsurface can have non-compact “crowned boundaries.” This section collects results about hyperbolic structures on such “crowned surfaces” as well as the relationship between properties of the lamination and the topology of its complementary subsurfaces.

Remark 4.1. Throughout this section and the following, we reserve S to denote a closed surface. If λ is a geodesic lamination, then $S \setminus \lambda$ denotes the metric completion of the complementary subsurfaces to λ (with respect to some auxiliary hyperbolic metric); we will refer to the topological type of a component of $S \setminus \lambda$ by Σ . Hyperbolic metrics on S and Σ will be denoted by X and Y , respectively.

Hyperbolic crowns. While less familiar than surfaces with boundary, crowned hyperbolic surfaces naturally arise by uniformizing surfaces with boundary and marked points on the boundary. They are also intricately related to meromorphic differentials on Riemann surfaces with high order poles (see, e.g., [Gup17]).

A *hyperbolic crown* with c_k spikes is a complete, finite-area hyperbolic surface with geodesic boundary that is homeomorphic to an annulus with c_k points removed from one boundary component. In the hyperbolic metric, the circular boundary component corresponds to a closed geodesic and each interval of the other boundary becomes a bi-infinite geodesic running between ideal vertices; compare Figure 2.

In general, a *hyperbolic surface with crowned boundary* is a complete, finite-area hyperbolic surface with totally geodesic boundary; the boundary components are either compact or hyperbolic crowns. We record the topological type of a crowned surface of genus g with b closed boundary components and k crowns with c_1, \dots, c_k many spikes as $\Sigma_{g,b}^{\{\underline{c}\}}$, where $\{\underline{c}\} = \{c_1, \dots, c_k\}$.

Remark 4.2. Ideal polygons may be considered as crowned surfaces of genus 0 with a single (crowned) boundary component. All of the results in this section hold for both crowned surfaces with nontrivial topology as well as for ideal polygons, but their proofs are slightly different. Our citations of [Gup17] are all for the case when Σ is not an ideal polygon; for the corresponding statements for ideal polygons, see [Gup17, Section 3.3] or [HTTW95].

Every crowned surface Y with non-cyclic (and non-trivial) fundamental group contains a “convex core” obtained by cutting off its crowns along a geodesic multicurve [CB88, Lemma 4.4]. When Y has type $\Sigma_{g,b}^{\{\underline{c}\}}$, this core is a subsurface of genus g with $b+k$ closed boundary components. Since each crown with c_i spikes may be decomposed into c_i ideal hyperbolic triangles by introducing leaves wrapping around the totally

geodesic boundary component, we have the following expression for the area:

$$(4) \quad \frac{1}{\pi} \text{Area}(Y) = 4g - 4 + 2b + \sum_{i=1}^k (c_i + 2).$$

Note that one can triangulate an ideal polygon of c sides into $(c-2)$ ideal triangles, and so the above formula also holds for ideal polygons.

The metric residue. While crown ends (and ideal polygons) do not have well-defined boundary lengths, one can define a natural generalization when there are an even number of spikes. This turns out to be a fundamental invariant that controls when crowns can be glued together along a lamination (Lemma 13.1).

Let \mathcal{C} be a hyperbolic crown or an ideal polygon with c spikes, where c is even. One can then orient \mathcal{C} , that is, pick an orientation of the boundary leaves so that the orientations of asymptotic leaves agree. Truncating each spike of \mathcal{C} along a horocycle based at the tip of the spike yields a surface with a boundary made up of horocyclic segments h_1, \dots, h_c and geodesic segments g_1, \dots, g_c . See Figure 2.

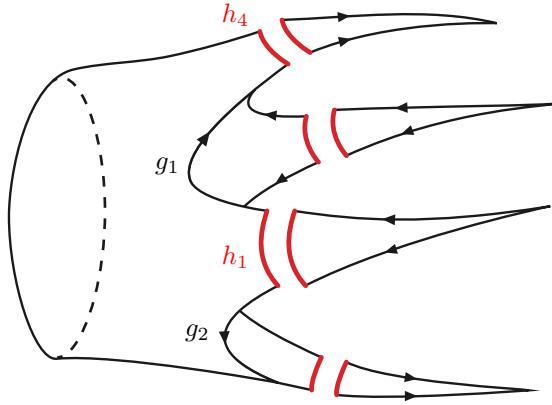


FIGURE 2. Truncating an (oriented) crown to compute its metric residue.

Definition 4.3 (Definition 2.9 of [Gup17]). Let \mathcal{C} be either an oriented hyperbolic crown or an oriented ideal polygon with an even number of spikes. Then its *metric residue* $\text{res}(\mathcal{C})$ is

$$\text{res}(\mathcal{C}) = \sum_{i=1}^c \varepsilon_i \ell(g_i)$$

where ε_i is positive if the truncated crown lies on the left of g_i and negative if it lies on the right.

Since changing the truncation depth of a spike increases the length of two adjacent sides, the metric residue evidently does not depend on the choice of truncation [Gup17, Lemma 2.10]. Observe also that flipping the orientation of \mathcal{C} flips the sign of its metric residue.

Similarly, define the metric residue of an oriented totally geodesic boundary component β of Y to be $\pm \ell(\beta)$, where the sign depends on whether Y lies to the left of β (positive) or right (negative).

Deformation spaces of crowned surfaces. We now record some useful facts about the Teichmüller spaces of crowned hyperbolic surfaces.

Given any crowned hyperbolic surface Y , one can obtain a natural compactification \widehat{Y} by adding on an ideal vertex at the end of each spike of each crown. The corresponding (topological) surface $\widehat{\Sigma}_{g,b}^{\{c\}}$ then has $b+k$ boundary components with c_i marked points on the $(b+i)^{\text{th}}$ boundary component. A *marking* of a crowned hyperbolic surface Y is a homeomorphism

$$f : \widehat{\Sigma}_{g,b}^{\{c\}} \rightarrow \widehat{Y}$$

which takes boundary marked points to ideal vertices. We think of the boundary marked points as having distinct labels, so different identifications of the boundary points of $\widehat{\Sigma}_{g,b}^{\{c\}}$ with the spikes of Y yield different markings. The Teichmüller space of a crowned hyperbolic surface $\Sigma_{g,b}^{\{c\}}$ is then defined to be the space of

all marked hyperbolic metrics on $\Sigma_{g,b}^{\{\mathcal{C}\}}$, up to isotopies which fix the totally geodesic boundary components pointwise and fix each ideal vertex of each crown.

As noted above, any crowned hyperbolic surface $\Sigma_{g,b}^{\{\mathcal{C}\}}$ contains an uncrowned subsurface which serves as its convex core. Therefore, the Teichmüller space of a crowned hyperbolic surface may be parametrized by the Teichmüller space of its convex core together with parameters describing each crown and how it is attached. A precise version of this dimension count is recorded below.

Lemma 4.4 (Lemma 2.16 of [Gup17]). Let $\Sigma = \Sigma_{g,b}^{\{\mathcal{C}\}}$ be a crowned hyperbolic surface or an ideal polygon. Then $\mathcal{T}(\Sigma) \cong \mathbb{R}^d$, where

$$(5) \quad d = 6g - 6 + 3b + \sum_{i=1}^k (c_i + 3).$$

Fixing the length of any closed boundary component of $\Sigma_{g,b}^{\{\mathcal{C}\}}$ cuts out a codimension 1 subvariety of $\mathcal{T}(\Sigma)$. Similarly, the subspace of surfaces with fixed metric residues at an even-spiked crown has codimension one. The following proposition ensures that the intersections of the level sets of length and metric residue are topologically just cells of the proper dimension:

Proposition 4.5 (Corollary 2.17 in [Gup17]). Let $\Sigma = \Sigma_{g,b}^{\{\mathcal{C}\}}$ be a crowned surface or an ideal polygon. Let β_1, \dots, β_b denote the closed boundary components of Σ and let $\mathcal{C}_1, \dots, \mathcal{C}_e$ denote the crown ends which have an even number of spikes. Fix an orientation of each crown end. Then for any $(L_i) \in \mathbb{R}_{>0}^b$ and any $(R_j) \in \mathbb{R}^e$,

$$\{(Y, f) \in \mathcal{T}(\Sigma) \mid \ell(\beta_i) = L_i \text{ and } \text{res}(\mathcal{C}_j) = R_j \text{ for all } i, j\} \cong \mathbb{R}^{d-b-e}$$

where d is as in (5).

Topology. When a crowned surface Σ comes from cutting a closed surface S along a geodesic lamination λ , we can relate the topology of λ to the topology of Σ .

Recall that the Euler characteristic of a lamination is defined to be alternating sum of the ranks of its Čech cohomology groups, viewing λ as a subset of S . Below, we compute the Euler characteristic of a geodesic lamination in terms of the topological type of its complementary subsurfaces.

Lemma 4.6. Let λ be a geodesic lamination on S . Then the total number of spikes of $S \setminus \lambda$ equals $-2\chi(\lambda)$.

We also record the corresponding formula for later use. Suppose that $\overline{S \setminus \lambda} = \Sigma_1 \cup \dots \cup \Sigma_m$; then

$$(6) \quad \chi(\lambda) = -\frac{1}{2} \sum_{j=1}^m \sum_{i=1}^{k_j} c_i^j$$

where $\{c_1^j, \dots, c_{k_j}^j\}$ denotes the crown type of Σ_j .

Proof. Fix some train track τ which carries λ and has the same topological type; in Section 5.2 below, this is referred to as *snug* carrying of λ on τ . Lemma 13 of [Bon97b] states that for any such train track, $\chi(\lambda) = \chi(\tau)$, and so it suffices to compute the Euler characteristic of τ .

Splitting the switches of τ if necessary, we may assume that τ is trivalent (observe that this operation preserves the Euler characteristic). Then each spike of $S \setminus \lambda$ corresponds to a unique switch of τ , and each switch corresponds to three half-edges, so

$$\# \text{ spikes}(S \setminus \lambda) = \# \text{ switches}(\tau) = \frac{2}{3} \cdot \# \text{ edges}(\tau).$$

Plugging this into the formula $\chi(\tau) = \# \text{ switches}(\tau) - \# \text{ edges}(\tau)$ proves the claim. \square

In general, the relationship between the boundary components of $S \setminus \lambda$ and λ can be rather involved. For example, one can construct a lamination on a closed surface of genus $g \geq 2$ consisting of 3 leaves, two of which are non-isotopic simple closed curves and one leaf which spirals onto each of the closed leaves. In this scenario, there is not a precise correspondence between closed leaves of λ and totally geodesic boundary components of its complementary subsurface.

When λ also supports a measure of full support, however, things become much nicer. In particular, each component of λ is minimal (in that every leaf is dense in the component) and so the closed leaves of λ are

all isolated. In this case, there is a natural 1-to-2 correspondence between closed leaves of λ and totally geodesic boundary components of $S \setminus \lambda$.

N.B. So that we do not have to deal with possible spiraling behavior of λ , we henceforth restrict our discussion to those laminations that support a measure. See also Remark 7.10.

5. THE ORTHOGEODESIC FOLIATION

In this section we construct the *orthogeodesic foliation* $\mathcal{O}_\lambda(X) \in \mathcal{MF}(\lambda)$ of a hyperbolic surface X with respect to λ and describe some of its basic properties.

5.1. The spine of a hyperbolic surface. We begin by describing the orthogeodesic foliation restricted to subsurfaces Y complementary to λ . Let Y be a finite area hyperbolic surface with totally geodesic boundary, possibly with crowned boundary. As we are most interested in the Y coming from cutting a closed surface along a lamination, we also assume that Y has no annular cusps.

Definition 5.1. The orthogeodesic foliation $\mathcal{O}_{\partial Y}(Y)$ of Y is the (singular, piecewise-geodesic) foliation of Y whose leaves are fibers of the closest point projection to ∂Y .

Near ∂Y , the leaves of $\mathcal{O}_{\partial Y}(Y)$ are geodesic arcs meeting ∂Y orthogonally. To understand the global structure of the foliation, however, we need to determine how the leaves extend into the interior of Y . In particular, we must understand the locus of points that are closest to multiple points of ∂Y .

To that end, for any point $x \in Y$, define the *valence* of x to be

$$\text{val}(x) := \#\{y \in \partial Y : d(x, y) = d(x, \partial Y)\}.$$

The (*geometric*) *spine* $\text{Sp}(Y)$ of Y is the set of points of Y with valence at least 2, and has a natural partition into subsets $\text{Sp}_k(Y)$, where $x \in \text{Sp}_k(Y)$ if it is equidistant from exactly k points in ∂Y . For the rest of the section we fix a hyperbolic surface Y and refer to $\text{Sp}(Y)$ and $\text{Sp}_k(Y)$ simply as Sp and Sp_k .

It is not hard to see that Sp is a properly embedded, piecewise geodesic 1-complex with some nodes of valence 1 removed (equivalently, a ribbon graph with some half-infinite edges). Indeed, Sp decomposes into a *finite core* Sp^0 and a finite collection of open geodesic rays; since we assumed Y had no annular cusps, each ray corresponds with a spike of a crowned boundary component. See [Mon09b, Section 2] for a discussion of the structure of the spine of a compact hyperbolic surface with geodesic boundary in which $\text{Sp}^0 = \text{Sp}$.

We record below a summary of this discussion; see also Figure 3.

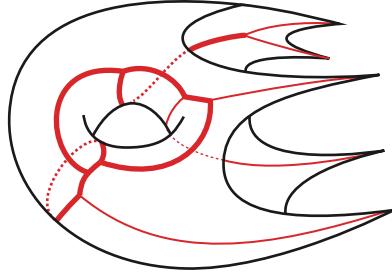


FIGURE 3. The spine of a hyperbolic surface with crowned boundary. Note that the finite core Sp^0 (represented in bold) contains a spine for the convex core of the surface.

Lemma 5.2. The finite core Sp^0 is a piecewise geodesically embedded graph, whose edges correspond to the components of Sp_2 with finite hyperbolic length and vertex set $\cup_{k \geq 3} \text{Sp}_k$. Each geodesic ray of $\text{Sp} \setminus \text{Sp}^0$ exits a unique spike of Y .

By definition, the orthogeodesic foliation $\mathcal{O}_{\partial Y}(Y)$ has k -pronged singular leaves emanating from $\cup_{k \geq 3} \text{Sp}_k$ for $k \geq 3$. The nonsingular leaves of $\mathcal{O}_{\partial Y}(Y)$ glue along $\text{Sp}_2(Y)$ (usually at an angle) and can be smoothed by an arbitrarily small isotopy supported near Sp . As the geometry of Sp interacts nicely with the leaves of $\mathcal{O}_{\partial Y}(Y)$, we generally prefer to think about $\mathcal{O}_{\partial Y}(Y)$ as a piecewise geodesic singular foliation rather than as a smooth one. When convenient, we will pass freely between the orthogeodesic foliation and a smoothing.

We observe that there is also an isotopy supported in the ends of the spikes of Y and restricting to the identity on ∂Y that maps leaves of the orthogeodesic foliation to horocycles based at the tip of the spike. This equivalence between the orthogeodesic and horocyclic foliations in spikes is of vital importance in Sections 13 [15] as it allows us to adapt many of Bonahon and Thurston's arguments to this setting.

Remark 5.3. One can check that, for regular ideal polygons, the isotopy in spikes extends to a global isotopy between the orthogeodesic foliation and the symmetric partial foliation by horocycles.

Following the leaves of the orthogeodesic foliation in the direction of \mathbf{Sp} defines a deformation retraction of Y onto \mathbf{Sp} ; let $r : Y \rightarrow \mathbf{Sp}$ be the map fully collapsing Y onto \mathbf{Sp} . For x and y in the same component of \mathbf{Sp}_2^0 , the leaves $r^{-1}(x)$ and $r^{-1}(y)$ of $\mathcal{O}_{\partial Y}(Y)$ are properly isotopic. We may therefore associate to each edge e of \mathbf{Sp}_2 the (proper) isotopy class of $r^{-1}(x)$ for $x \in e$; we call this the *dual arc* α_e to e .

There is a distinguished representative of α_e that is geodesic and orthogonal to both ∂Y and e ; compare Figure 7. By abuse of notation, we henceforth identify α_e with its orthogeodesic representative and define

$$\underline{\alpha}(Y) := \bigcup_{e \subset \mathbf{Sp}_2^0} \alpha_e.$$

Lemma 5.4. The metric completion of the surface with corners $Y \setminus \underline{\alpha}(Y)$ is homeomorphic to a union of closed disks and closed disks with finitely many points on the boundary removed. That is, $\underline{\alpha}(Y)$ fills Y .

Proof. Each component of $Y \setminus \underline{\alpha}(Y)$ deformation retracts onto a component of the metric completion of $\mathbf{Sp} \setminus \underline{\alpha}(Y)$. By the duality of arcs and edges of \mathbf{Sp}_2^0 , each component of $\mathbf{Sp} \setminus \underline{\alpha}(Y)$ is contractible. \square

The orthogeodesic foliation also comes with a natural transverse measure: the measure of an arc k transverse to (a smoothing of) $\mathcal{O}_{\partial Y}(Y)$ is defined on small enough transverse arcs k first by isotoping the arc into ∂Y transversely to $\mathcal{O}_{\partial Y}(Y)$ and then measuring the hyperbolic length there. Locally, the orthogeodesic foliation admits a reflection about each edge of \mathbf{Sp} , so by restricting k to those leaves of $\mathcal{O}_{\partial Y}(Y)$ that intersect a given edge, we can use this symmetry to see that the measure of k is the same after a transverse isotopy onto either boundary component of Y . Extending to all transverse arcs by additivity defines a transverse measure on $\mathcal{O}_{\partial Y}(Y)$.

To each component e of \mathbf{Sp}_2^0 we associate the length $c_e > 0$ of either component of $r^{-1}(e) \cap \partial Y$; the transverse measure of e is exactly c_e . Anticipating the contents of the next section (see, e.g., Theorem 6.4), we define the formal sum

$$(7) \quad \underline{A}(Y) := \sum_{e \subset \mathbf{Sp}_2^0} c_e \alpha_e.$$

5.2. The orthogeodesic foliation. Now that we have described the orthogeodesic foliation on each component of $S \setminus \lambda$, we can glue these pieces together along the leaves of λ to get a foliation of S .

Construction 5.5. Let $X \in \mathcal{T}(S)$ and λ be a geodesic lamination on X . Cutting X open along λ taking the metric completion of each component, we obtain a union of hyperbolic surfaces with totally geodesic boundary (possibly with crowned boundary). On each such component Y , we construct the orthogeodesic foliation $\mathcal{O}_{\partial Y}(Y)$ as described in Section 5.1 above.

A standard fact from hyperbolic geometry [CEG06, Lemma 5.2.6] shows that the line field defined by (a smoothing of) the orthogeodesic foliation forms a Lipschitz line field on $X \setminus \lambda$. Since λ has measure 0, this line field is integrable near λ , so the partial foliation defined on $X \setminus \lambda$ extends across the leaves of λ . This defines a measured foliation $\mathcal{O}_\lambda(X) \in \mathcal{MF}(S)$, hence a map $\mathcal{O}_\lambda : \mathcal{T}(S) \rightarrow \mathcal{MF}(S)$.

Later, we prove Lemma 5.8 that λ and $\mathcal{O}_\lambda(X)$ bind, allowing us to restrict the codomain of \mathcal{O}_λ to $\mathcal{MF}(\lambda)$. Ultimately, our goal is to show that \mathcal{O}_λ is a homeomorphism onto $\mathcal{MF}(\lambda)$.

Geometric train tracks. We now consider the geometry of $\mathcal{O}_\lambda(X)$ in a neighborhood of λ . The following is a modification of an important construction of Thurston [Thu82, Chapter 8.9].

Construction 5.6. Let $\epsilon > 0$ be small enough so that the ϵ -neighborhood $\mathcal{N}_\epsilon(\lambda)$ is topologically stable. The orthogeodesic foliation $\mathcal{O}_\lambda(X)$ restricts to a foliation of $\mathcal{N}_\epsilon(\lambda)$ without singular points, and collapsing the leaves yields a quotient map $\pi : \mathcal{N}_\epsilon(\lambda) \rightarrow \tau$ where τ can be C^1 -embedded in $\mathcal{N}_\epsilon(\lambda)$ as a train track

carrying λ in X . By changing ϵ , we may assume that τ is trivalent. ⁶ Then $\tau = \tau(\lambda, X, \epsilon)$ is a *geometric train track*.

We sometimes refer to $\mathcal{N}_\epsilon(\lambda)$ as a *train track neighborhood* of λ and the leaves of $\mathcal{O}_\lambda(X)|_{\mathcal{N}_\epsilon(\lambda)}$ as *ties*. A train track neighborhood coming from Construction 5.6 is a union of bands and annuli foliated by ties glued together along the ties that collapse to switches of τ . We recall that if λ meets every tie of τ and there is no path between spikes of $S \setminus \mathcal{N}_\epsilon(\lambda)$ that is contained in $\mathcal{N}_\epsilon(\lambda) \setminus \lambda$, then τ is said to *snugly carry* λ . Equivalently, τ snugly carries λ if and only if $S \setminus \lambda$ and $S \setminus \tau$ have the same topological type. With this definition, it is clear that the geometric train tracks constructed above always carry λ snugly.

Using the geometry of $\pi : \mathcal{N}_\epsilon(\lambda) \rightarrow \tau$, the branches of τ admit a well defined notion of length. Indeed, let $b \subset \tau$ be a branch, and choose a lift \tilde{b} to the universal cover \tilde{X} . Let $\ell, \ell' \subset \tilde{\lambda}$ be leaves of the elevation $\tilde{\lambda}$ of λ to \tilde{X} that meet $\pi^{-1}(\tilde{b}) \subset \mathcal{N}_\epsilon(\tilde{\lambda})$ in segments g and g' . Since $\mathcal{O}_\lambda(X)$ is equivalent to a horocyclic foliation in $\mathcal{N}_\epsilon(\lambda)$, transporting g along the leaves of $\mathcal{O}_{\tilde{\lambda}}(\tilde{X})$ near \tilde{b} onto g' is isometric, so $\ell_X(g) = \ell_X(g')$. We may therefore define the *length* of b (along λ) as

$$\ell_X(b) := \ell_X(g)$$

for any g as above. Similarly, for any branch $b \subset \tau$, the ties of $\mathcal{N}_\epsilon(\lambda)$ collapsing to b all have the same integral with respect to λ . Define

$$\lambda(b) := \lambda(k)$$

for any tie $k \subset \mathcal{O}_\lambda(X)|_{\pi^{-1}(b)}$; this is equivalently the weight deposited by λ on b in its τ train track coordinates.

Lemma 5.7. For any hyperbolic structure X and any measure λ' on λ , we have $i(\lambda', \mathcal{O}_\lambda(X)) = \ell_X(\lambda')$.

Proof. Using Construction 5.6, find a geometric train track $\pi : \mathcal{N}_\epsilon(\lambda) \rightarrow \tau$ snugly carrying λ on X . By definition, the intersection pairing is given by the integral over X of the product measure $d\lambda' \otimes d\mathcal{O}_\lambda(X)$, whose support is contained entirely in the train track neighborhood $\mathcal{N}_\epsilon(\lambda)$. For each branch $b \subset \tau$, the integral of this measure on $\pi^{-1}(b)$ is just $\lambda'(b)\ell_X(b)$, so

$$\begin{aligned} i(\lambda', \mathcal{O}_\lambda(X)) &= \iint_X d\lambda' \otimes d\mathcal{O}_\lambda(X) \\ &= \iint_{\mathcal{N}_\epsilon(\lambda)} d\lambda' \otimes d\mathcal{O}_\lambda(X) = \sum_{b \subset \tau} \iint_{\pi^{-1}(b)} d\lambda' \otimes d\mathcal{O}_\lambda(X) \\ &= \sum_{b \subset \tau} \lambda'(b)\ell_X(b). \end{aligned}$$

On the other hand, $\ell_X(\lambda')$ is the integral over X of the measure $d\lambda' \otimes dl_{\lambda'}$, locally the product of the transverse measure λ' and 1-dimensional Lebesgue measure $l_{\lambda'}$ on the support of λ' . Since λ' is supported in λ , the integral of $d\lambda' \otimes dl_{\lambda'}$ is equal to the integral of $d\lambda' \otimes dl_\lambda$, and again the support of the product measure is contained in $\mathcal{N}_\epsilon(\lambda)$. On each thickened branch $\pi^{-1}(b) \subset \mathcal{N}_\epsilon(\lambda)$, the integral of $d\lambda' \otimes dl_\lambda$ is $\lambda'(b)\ell_X(b)$, giving the equality

$$\ell_X(\lambda') = \sum_{b \subset \tau} \lambda'(b)\ell_X(b).$$

This completes the proof of the lemma. \square

With this computation, we can now show that λ and $\mathcal{O}_\lambda(X)$ together bind S .

Lemma 5.8. For any $X \in \mathcal{T}(S)$ and $\lambda \in \mathcal{ML}(S)$, we have $\mathcal{O}_\lambda(X) \in \mathcal{MF}(\lambda)$.

Proof. Suppose that η is a measured lamination so that $i(\eta, \lambda) = 0$; without loss of generality, we may assume that η is ergodic. Then one of two things must be true: either η is supported on λ or its support is disjoint from λ . In the first case, $i(\eta, \mathcal{O}_\lambda(X)) = \ell_X(\eta) > 0$ by Lemma 5.7.

If η is disjoint from λ then it is contained in a component Y of $X \setminus \lambda$, and we need only show that $i(\eta, \mathcal{O}_\lambda(X)) > 0$. Scaling the measure of η as necessary, let us assume that $\ell_X(\eta) = \ell_Y(\eta) = 1$. Now we recall that the set of weighted simple closed curves is dense in the space of measured laminations on Y . By

⁶In the literature, trivalent train tracks are also called “generic.”

homogeneity and continuity of the intersection pairing, it therefore suffices to find some uniform $\epsilon > 0$ such that

$$i(\gamma, \mathcal{O}_\lambda(X)) \geq \epsilon \ell_X(\gamma)$$

for every simple closed curve $\gamma \subset Y$. Indeed, once we have demonstrated such a bound we may approximate η arbitrarily well by weighted curves $\gamma/\ell_X(\gamma)$ to deduce the desired bound on $i(\eta, \mathcal{O}_\lambda(X))$.

So let Y_0 be the convex hull of $r^{-1}(\mathbf{Sp}^0)$; Y_0 is compact and the inclusion of Y_0 into Y is a homotopy equivalence. Any simple closed geodesic γ in Y is contained in Y_0 , and since Y deformation retracts onto the component of \mathbf{Sp} contained in Y , γ is homotopic to a concatenation of edges in \mathbf{Sp}^0 .

Give \mathbf{Sp}^0 a metric making its edges e have length $c_e = i(e, \mathcal{O}_\lambda(X))$; then the inclusion $\mathbf{Sp}^0 \rightarrow Y_0$ with this metric induces an equivariant quasi-isometry on universal covers (this follows because they are both Gromov hyperbolic and $\pi_1(Y)$ acts cocompactly and properly discontinuous on each). The geodesic lengths of closed curves in \mathbf{Sp} and in Y_0 are therefore comparable, so that there is some $\epsilon > 0$ so that

$$i(\gamma, \mathcal{O}_\lambda(X)) = \ell_{\mathbf{Sp}^0}(\gamma) \geq \ell_X(\gamma)\epsilon,$$

demonstrating the desired uniform bound. \square

5.3. Deflation. For a given pair $(X, \lambda) \in \mathcal{T}(S) \times \mathcal{ML}(S)$, the pair of laminations $\mathcal{O}_\lambda(X)$ and λ bind by Lemma 5.8. By Theorem 2.1, there is a unique quadratic differential $q = q(\mathcal{O}_\lambda(X), \lambda)$, holomorphic on some Riemann surface Z , whose real and imaginary foliations are $\mathcal{O}_\lambda(X)$ and λ , respectively. In this section we define a *deflation* map $\mathcal{D}_\lambda : X \rightarrow Z$ that allows us to make direct comparisons between the hyperbolic geometry of X and the singular flat geometry q .

An informal description of \mathcal{D}_λ is that it “deflates” the subsurfaces of $X \setminus \lambda$, retracting them to \mathbf{Sp} along the leaves of $\mathcal{O}_\lambda(X)$, while it “inflates” along the leaves of λ according to the transverse measure. The orthogeodesic foliation in a neighborhood of λ assembles into the vertical foliation of the resulting quadratic differential metric and \mathcal{D}_λ maps $\mathbf{Sp} \subset X$ to the horizontal separatrices; compare Figure 4.

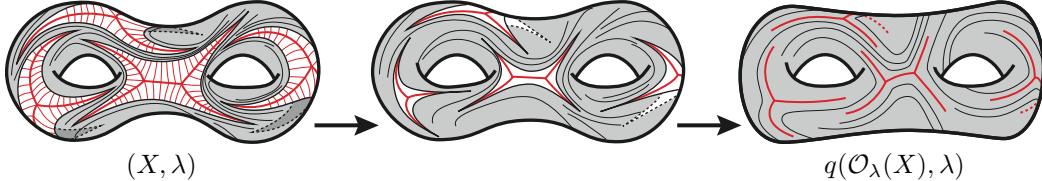


FIGURE 4. Inflating a lamination and deflating its complementary components.

Remark 5.9. This heuristic description of \mathcal{D}_λ can be made precise by grafting X along λ (see, e.g., [Dum09]) and then collapsing the hyperbolic pieces along the leaves of $\mathcal{O}_\lambda(X)$. In particular, \mathcal{D}_λ is *not* the grafting map.

Proposition 5.10. Given a marked hyperbolic structure $[f : S \rightarrow X] \in \mathcal{T}(S)$ and $\lambda \in \mathcal{ML}(S)$, let $[g : S \rightarrow Z] \in \mathcal{T}(S)$ be the marked complex structure on which $q(\mathcal{O}_\lambda(X), \lambda)$ is holomorphic. There is a map

$$\mathcal{D}_\lambda : X \rightarrow Z$$

that is a homotopy equivalence restricting to an isometry between \mathbf{Sp}^0 with its metric induced by integrating the edges against $\mathcal{O}_\lambda(X)$ and the graph of horizontal saddle connections of $q(\mathcal{O}_\lambda(X), \lambda)$ with the induced path metric. Moreover, $\mathcal{D}_\lambda \circ f \sim g$ and $\mathcal{D}_{\lambda*} \mathcal{O}_\lambda(X) = \text{Re}(q)$ and $\mathcal{D}_{\lambda*} \lambda = \text{Im}(q)$ as measured foliations.

Proof. Construction 5.6 supplies us with a geometric train track $\pi : \mathcal{N}_\epsilon(\lambda) \rightarrow \tau$. On the preimage $\pi^{-1}(b)$ of each closed branch b of τ we integrate the two measures $\mathcal{O}_\lambda(X)|_{\mathcal{N}_\epsilon(\lambda)}$ and λ giving $\pi^{-1}(b)$ the structure of a bi-foliated Euclidean rectangle of length $\ell_X(b)$ and height $\lambda(b)$. These rectangles glue along their ‘short’ sides $\{\pi^{-1}(s) : s \text{ is a switch of } \tau\}$ to give $\mathcal{N}_\epsilon(\lambda)$ the structure of a bi-foliated Euclidean band complex.

The map π extends to a self-homotopy equivalence of X homotopic to the identity preserving the orthogeodesic foliation leafwise. This means that the boundary of $\mathcal{N}_\epsilon(\lambda)$ admits a natural retraction onto \mathbf{Sp} by

⁷Throughout the paper we suppress markings in our notation, but reintroduce them here to state the proposition precisely.

collapsing the leaves of the orthogeodesic foliation in the complement of $\mathcal{N}_\epsilon(\lambda)$, and we take the quotient generated by this equivalence relation to obtain a new surface Y with its complex structure described below.

On each rectangle $\pi^{-1}(b)$, the bi-foliated Euclidean structure gives local coordinates to \mathbb{C} away from the singular points of $\mathcal{O}_\lambda(X)$ locally mapping $\mathcal{O}_\lambda(X)$ to $|dx|$ and λ to $|dy|$, thought of as measured foliations on the plane. These coordinate patches glue together along the spine to give local coordinates away from the points of valence ≥ 3 . Moreover, these charts preserve $|dx|$ and $|dy|$, so the transitions must be of the form $z \mapsto \pm z + \alpha$ for some $\alpha \in \mathbb{C}$. We have therefore built a Riemann surface Z equipped with a half-translation structure away from the vertices of Sp , which become cone points of cone angle equal to $\pi \cdot \text{val}(v)$. Edges of Sp join vertices along horizontal trajectories representing all horizontal saddle connections on q ; their lengths in the singular flat metric are given by the integral over $\mathcal{O}_\lambda(X)$. Thus \mathcal{D}_λ induces an isometry of metric graphs, as claimed. \square

This explicit description of the quadratic differential associated to the pair (X, λ) by the map \mathcal{O} from the introduction will be useful in order to prove in Theorem 13.13 that Diagram (2) commutes.

6. CELLULATING CROWNED TEICHMÜLLER SPACES

We now define a certain arc complex which combinatorializes the structure the orthogeodesic foliation on complementary subsurfaces. The main result of this section is Theorem 6.4, which shows that this arc complex is equivariantly homeomorphic to the Teichmüller space of the complementary surface. In particular, this shows that the restriction of the orthogeodesic foliation to each component of $S \setminus \lambda$ completely determines the hyperbolic structure on that piece.

Before stating the theorem, we must first set up our combinatorial analogue for Teichmüller space. This appears as Definition 6.1 after a series of auxiliary constructions.

Suppose that $\Sigma = \Sigma_{g,b}^{\{c\}}$ is a finite-area hyperbolic surface with boundary and without annular cusps. A properly embedded arc $I \rightarrow \Sigma$ is *essential* if I cannot be isotoped (through properly embedded arcs) into $\partial\Sigma$ or into a spike. The *arc complex* $\mathcal{A}(\Sigma, \partial\Sigma)$ of Σ *rel boundary* is the (simplicial, flag) complex whose vertices are isotopy classes of simple essential arcs of Σ . Vertices span a simplex in $\mathcal{A}(\Sigma, \partial\Sigma)$ if and only if there exist a collection of pairwise disjoint representatives for each isotopy class. The *filling arc complex* $\mathcal{A}_{\text{fill}}(\Sigma, \partial\Sigma)$ is the subset of $\mathcal{A}(\Sigma, \partial\Sigma)$ consisting only of those arc systems which cut Σ into a union of topological disks.

The geometric realization $|\mathcal{A}(\Sigma, \partial\Sigma)|$ of $\mathcal{A}(\Sigma, \partial\Sigma)$ is obtained by declaring every simplex to be a regular Euclidean simplex of the proper dimension; note that the topology of $|\mathcal{A}(\Sigma, \partial\Sigma)|$ obtained from the metric structure is in general different from the standard simplicial topology (see, e.g., [BE88]). The geometric realization $|\mathcal{A}_{\text{fill}}(\Sigma, \partial\Sigma)|$ is then the subspace of filling arc systems equipped with the subspace topology induced by the metric structure.

Definition 6.1. The *weighted filling arc complex* $|\mathcal{A}_{\text{fill}}(\Sigma, \partial\Sigma)|_{\mathbb{R}}$ of Σ *rel boundary* is the set of all weighted multi-arcs of the form

$$\underline{A} = \sum c_i \alpha_i$$

where $\underline{\alpha} = \bigcup \alpha_i \in \mathcal{A}_{\text{fill}}(\Sigma, \partial\Sigma)$ and $c_i > 0$ for all i .

Throughout, we will use α to denote a single arc, and $\underline{\alpha}$ to denote an (unweighted) multi-arc. The symbol \underline{A} will be reserved to denote a weighted multi-arc.

Note. If Σ is an ideal hyperbolic polygon, then the empty arc system fills Σ and we consider it as an element of $|\mathcal{A}_{\text{fill}}(\Sigma, \partial\Sigma)|_{\mathbb{R}}$. If Σ is not a polygon, then the empty arc system never fills.

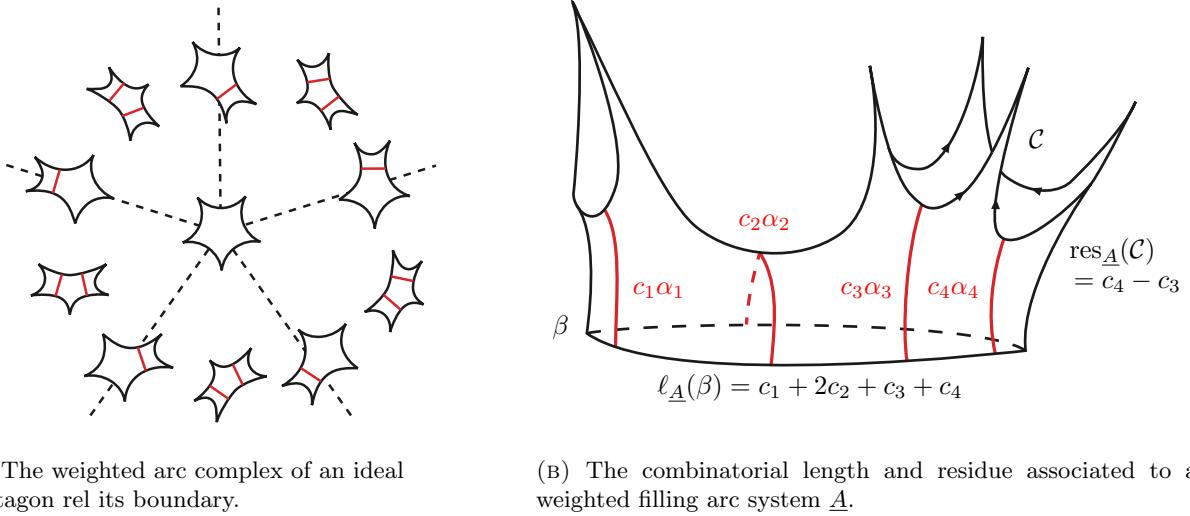
So long as Σ is not an ideal polygon, $|\mathcal{A}_{\text{fill}}(\Sigma, \partial\Sigma)|_{\mathbb{R}}$ is just $|\mathcal{A}_{\text{fill}}(\Sigma, \partial\Sigma)| \times \mathbb{R}_{>0}$. When Σ is an ideal polygon, then $|\mathcal{A}_{\text{fill}}(\Sigma, \partial\Sigma)|_{\mathbb{R}}$ is homeomorphic to the open cone on the filling arc complex:

$$(|\mathcal{A}_{\text{fill}}(\Sigma, \partial\Sigma)| \times \mathbb{R}_{\geq 0}) / (|\mathcal{A}_{\text{fill}}(\Sigma, \partial\Sigma)| \times \{0\}).$$

See Figure 5a for an example of $|\mathcal{A}_{\text{fill}}(\Sigma, \partial\Sigma)|_{\mathbb{R}}$ in the case when Σ is an ideal pentagon.

Remark 6.2. The standard duality between arc systems and ribbon graphs (see, e.g., [Mon09a]) assigns to every $\underline{A} \in |\mathcal{A}_{\text{fill}}(\Sigma, \partial\Sigma)|_{\mathbb{R}}$ a metric ribbon graph spine for Σ (with some infinitely long edges if Σ has crowns). One could of course translate the cell structure of $|\mathcal{A}_{\text{fill}}(\Sigma, \partial\Sigma)|_{\mathbb{R}}$ into a cellulation of an appropriate space of marked metric ribbon graphs.

While the arc complex definition is more practical for our definition of shear-shape cocycles, the dual ribbon graph picture allows us to immediately understand how to record the geometry of the horizontal trajectories of a quadratic differential (see Section 10).



(A) The weighted arc complex of an ideal pentagon rel its boundary.

(B) The combinatorial length and residue associated to a weighted filling arc system \underline{A} .

FIGURE 5. Arc complexes and combinatorial geometry.

Combinatorial geometry. Now that we have defined our combinatorial analogue of Teichmüller space, we can also define combinatorial notions of both length and metric residue.

Suppose that β is a compact boundary component of Σ and $\underline{A} \in |\mathcal{A}_{\text{fill}}(\Sigma, \partial\Sigma)|_{\mathbb{R}}$; then we define the \underline{A} -length $\ell_{\underline{A}}(\beta)$ of β to be the sum of the weights of the arcs of \underline{A} incident to β (counted with multiplicity, so that if both endpoints of $\underline{\alpha}$ lie on β then its weight is counted twice).

Similarly, let C be an oriented crowned boundary component with an even number of spikes. Then the edges of C are partitioned into those that have the surface lying on their left and those which have the surface on their right; call these edges positively and negatively oriented, respectively. The \underline{A} -residue $\text{res}_{\underline{A}}(C)$ of C is then defined to be the sum of the weights of the arcs incident to each positively oriented edge of C minus the sum of the weights of the arcs incident to the negatively oriented edges (where both sums are again taken with multiplicity). See Figure 5b for an example calculation.

We have now come to the most important object of this section, and a foundational result of this paper that allows us to pass between hyperbolic metrics, orthogeodesic foliations, and metric graphs embedded in flat structures.

Construction 6.3. Let Y be a crowned hyperbolic surface. As discussed in Section 5.1, the orthogeodesic foliation determines a spine for Y together with a dual (filling) arc system $\underline{\alpha}(Y)$. Weighting each dual arc by integrating the measure induced by $\mathcal{O}_{\partial Y}(Y)$ over the corresponding edge of Sp (compare (7)) therefore defines a map

$$\underline{A} : \mathcal{T}(\Sigma) \rightarrow |\mathcal{A}_{\text{fill}}(\Sigma, \partial\Sigma)|_{\mathbb{R}}.$$

When Σ has compact boundary, [Luo07, Theorem 1.2 and Corollary 1.4] states that $\underline{A}(\cdot)$ is a $\text{Mod}(\Sigma)$ -equivariant stratified real-analytic homeomorphism; see also [Mon09b, Do08, Ush99]. Our aim is to generalize Luo's theorem to surfaces with crowned boundary. While the arguments of [Luo07] can probably be adapted to this setting, we prefer to use some elementary hyperbolic geometry to realize $|\mathcal{A}_{\text{fill}}(\Sigma, \partial\Sigma)|_{\mathbb{R}}$ as a subcomplex sitting “at infinity” of the weighted filling arc complex of a surface with compact boundary.

Theorem 6.4. Let Σ be a crowned hyperbolic surface. Then the map

$$\underline{A} : \mathcal{T}(\Sigma) \rightarrow |\mathcal{A}_{\text{fill}}(\Sigma, \partial\Sigma)|_{\mathbb{R}}$$

is a $\text{Mod}(\Sigma)$ -equivariant stratified real analytic homeomorphism. Moreover, let β_1, \dots, β_b denote the closed boundary components of Σ and $\mathcal{C}_1, \dots, \mathcal{C}_e$ the crown ends which have an even number of spikes. Fix an orientation of each \mathcal{C}_j . Then the map above identifies the level sets

$$\{(Y, f) \in \mathcal{T}(\Sigma) \mid \ell(\beta_i) = L_i, \text{res}(\mathcal{C}_j) = R_j\} \cong \{\underline{A} \in |\mathcal{A}_{\text{fill}}(\Sigma, \partial\Sigma)|_{\mathbb{R}} : \ell_{\underline{A}}(\beta_i) = L_i, \text{res}_{\underline{A}}(\mathcal{C}_j) = R_j\}$$

for any $(L_i) \in \mathbb{R}_{>0}^b$ and any $(R_j) \in \mathbb{R}^e$.

The remainder of this section is devoted to deducing Theorem 6.4 from [Luo07, Theorem 1.2, Corollary 1.4] and [Mon09b, Section 2.4]. Our plan is to appeal to the aforementioned references to prove that for a given maximal arc system $\underline{\alpha}$, the map $\underline{A}(\cdot)$ extends to a real analytic map $\underline{A}_{\underline{\alpha}} : \mathcal{T}(\Sigma) \rightarrow \mathbb{R}^{\underline{\alpha}}$ that agrees with $\underline{A}(\cdot)$ on the locus of hyperbolic surfaces whose spine has dual arc system contained in $\underline{\alpha}$ (Lemma 6.9). We show that $\underline{A}(\cdot)$ is a homeomorphism by building a continuous right inverse $Y : |\mathcal{A}_{\text{fill}}(\Sigma, \partial\Sigma)|_{\mathbb{R}} \rightarrow \mathcal{T}(\Sigma)$; $Y(\underline{A})$ is obtained as a geometric limit metrics on a larger compact surface with boundary as some arcs are pinched to spikes.

Endowing Σ with an auxiliary hyperbolic metric, we take Σ° to be the surface with geodesic and horocyclic boundary components obtained by truncating the tips of the spikes. Let $\underline{\gamma}$ be the union of horocyclic boundary components of Σ° and double Σ° along $\underline{\gamma}$ to obtain a (topological) surface $D\Sigma$ and an identification of Σ° with a subsurface of $D\Sigma$ taking $\partial\Sigma^\circ \setminus \underline{\gamma}$ into $\partial D\Sigma$; see Figure 6.

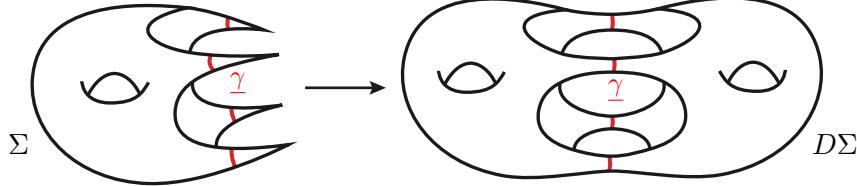


FIGURE 6. The truncation of a crowned surface Σ along $\underline{\gamma}$ and its double $D\Sigma$.

Let $\underline{A} = \sum c_i \alpha_i$ be a weighted filling arc system on Σ and let $\underline{\beta}$ be the mirror image of $\underline{\alpha}$ in $D\Sigma$, so that $\underline{\alpha} \cup \underline{\gamma} \cup \underline{\beta}$ is a filling arc system on $D\Sigma$. For each $t > 0$, define

$$\underline{B}_t = \sum c_i \beta_i + t \sum \gamma_i + \sum c_i \alpha_i \in |\mathcal{A}_{\text{fill}}(D\Sigma, \partial D\Sigma)|_{\mathbb{R}}.$$

Since $D\Sigma$ is compact, we can apply [Luo07, Corollary 1.4] which states that there is a unique hyperbolic structure $X_t \in \mathcal{T}(D\Sigma)$ whose natural weighted arc system coincides with \underline{B}_t .

Remark 6.5. It will be convenient to assume that $\underline{\alpha}$ is maximal, formally adding arcs of weight 0 to \underline{A} (and \underline{B}_t) as necessary.

Our goal is now to show that (X_t) converges as $t \rightarrow \infty$ to a surface $Y \in \mathcal{T}(\Sigma)$ such that $\underline{A}(Y) = \underline{A}$. The convergence is geometric: we take basepoints $x_t \in X_t$ lying outside of the “thin parts” of the subsurface corresponding to Σ° and extract a geometric limit of (X_t, x_t) as $t \rightarrow \infty$. The limit metric Y has spikes corresponding to $\underline{\gamma}$ and so defines a point in $\mathcal{T}(\Sigma)$. Moreover, Y inherits a filling arc system naturally identified with $\underline{\alpha}$, which is necessarily realized orthogeodesically.

We begin with an estimate on the lengths of orthogeodesic arcs.

Lemma 6.6. If X is a hyperbolic metric on a compact surface with totally geodesic boundary and $\underline{A}(X) = \sum c_i \alpha_i$, then

$$\min \left\{ \log 3, 2 \tanh^{-1} \left(\frac{\tanh(\log \sqrt{3})}{\cosh(c_i/2)} \right) \right\} \leq \ell_X(\alpha_i) \leq \frac{2\pi}{c_i},$$

for each i .

Proof. Any leaf of the orthogeodesic foliation properly homotopic to α_i has hyperbolic length at least $\ell_X(\alpha_i)$. Thus the embedded “collar” about α_i consisting of all leaves of the orthogeodesic foliation in the same

homotopy class of α_i has area at least $c_i \ell_X(\alpha_i)$ (see Figure 7). On the other hand, the Gauss–Bonnet theorem bounds area of the collar above by 2π , so we get a bound

$$\ell_X(\alpha_i) \leq \frac{2\pi}{c_i}.$$

Now we would like to find a lower bound for $\ell_X(\alpha_i)$ in terms of c_i ; for notational convenience we fix i and set $\alpha = \alpha_i$ and $c = c_i$. Assume that $\ell_X(\alpha) < \log 3$. Let H be a component of $X \setminus \underline{\alpha}$ meeting α ; then there is a unique point $u \in H$ equidistant from all boundary components of X meeting H . There is also a universal lower bound to the distance from u to any such boundary component, given by $\log \sqrt{3}$, the radius of the circle inscribed in an ideal triangle. Thus the leaf of $\mathcal{O}_{\partial X}(X)$ through u has length at least $\log(3)$.

Since $\ell_X(\alpha) < \log 3$, we know that there is a leaf of the orthogeodesic foliation parallel to α with length $\log 3$. Using a formula relating the lengths of the sides of a hyperbolic tri-rectangle [Bus10, Theorem 2.3.1], the distance c_0 from α and this leaf is given by

$$(8) \quad \tanh(\ell_X(\alpha)/2) = \frac{\tanh(\log \sqrt{3})}{\cosh(c_0)}.$$

Now this expression is decreasing in c_0 , and $x \mapsto \tanh^{-1}(x)$ is increasing. We have that $c > 2c_0$ by definition (see Figure 7), so the lemma follows. \square

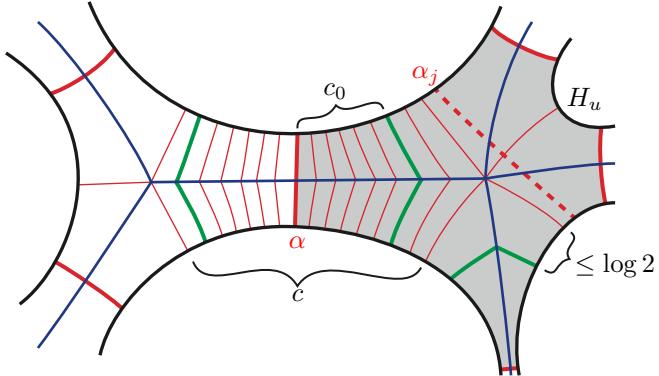


FIGURE 7. A foliated collar of width c about an orthogeodesic arc α . If the arc is shorter than $\log 3$, then there are (bold green) leaves of this collar of length equal to $\log 3$. For a very short arc γ_i , the distance between the longest leaf of its collar and the leaf of length $\log 3$ is at most $\log 2$. The dashed arc α_j has weight 0 and corresponds to one of two possible choices of maximal completion of $\underline{\alpha}$.

For any arc γ_i of $\underline{\gamma}$, some elementary estimates similar to those given in the proof of Lemma 6.6 (compare Equation (8)) give $\ell_t(\gamma_i) = O(e^{-t/2})$. If α_i appears in B_t with coefficient $c_i = 0$, then Lemma 6.6 provides a lower bound of $\log 3$ for the length $\ell_t(\alpha_i)$ of α_i on X_t . We also have the following upper bound:

Lemma 6.7. If $c_j = 0$ for some j , then for t large enough, we have

$$\log 3 \leq \ell_t(\alpha_j) \leq 2 \sum c_i + 8\pi \sum \frac{1}{c_i} + |\underline{\gamma}| \log 144.$$

Proof. We remove all arcs of $\underline{\alpha} \cup \underline{\gamma} \cup \underline{\beta}$ with positive weight from X_t and let H_t be (the metric completion of) the right-angled polygon component that contains α_j . Our strategy is to find a path of controlled length contained in ∂H_t joining the endpoints of α_j .

Notice that ∂H_t alternates between segments of ∂X_t and arcs of $\underline{\alpha} \cup \underline{\gamma} \cup \underline{\beta}$ with positive weight. From Lemma 6.6, the total length of segments coming from arcs of $\underline{\alpha} \cup \underline{\beta}$ is at most $8\pi \sum 1/c_i$, because each arc of $\underline{\alpha} \cup \underline{\beta}$ can appear at most two times on ∂H_t . Similarly, from the construction of our coordinate system, the total length of the segments coming from ∂X_t that correspond to collars of arcs in $\underline{\alpha} \cup \underline{\beta}$ is at most $2 \sum c_i$.

Suppose some arc γ_i of $\underline{\gamma}$ forms a segment of ∂H_t . The distance between the leaf of the orthogeodesic foliation parallel to γ_i with length $\log(3)$ and the singular, longest leaf parallel to γ_i has distance uniformly bounded above by $\log 2$ for large values of t (see Figure 7). Truncate H_t by removing the leaves of the orthogeodesic foliation parallel to γ_i with length at most $\log 3$ to obtain a new (non-convex) geodesic polygon H_t° . An application of the Collar Lemma [Bus10, Theorem 4.1.1] to the double DX_t along its boundary shows that α_j does not enter the region of H_t that we removed.

Each arc γ_i of $\underline{\gamma}$ contributed at most $2t + O(e^{-t/2})$ to the length of ∂H_t . However, after truncating, each γ_i contributes at most $2(\log 2 + \log 3 + \log 2) = \log 144$ to the length of ∂H_t° . Putting together all of our estimates completes the proof. \square

For each $\alpha_i \in \underline{\alpha}$ with positive coefficient c_i in \underline{B}_t , the orthogeodesic length $\ell_t(\alpha_i)$ of α_i on X_t is bounded above and below by the positive real numbers independent of t provided by Lemma 6.6. If $c_i = 0$ for some i , then Lemma 6.7 provides bounds on $\ell_t(\alpha_i)$ independent of t . Therefore, there exists a subsequence t_k tending to infinity such that $(\ell_{t_k}(\alpha_i))$ converges to a positive number ℓ_i for each i , while $\ell_t(\gamma_i) = O(e^{-t/2})$ for each $\gamma_i \in \underline{\gamma}$.

The metric completion of $X_{t_k} \setminus (\underline{\alpha} \cup \underline{\gamma} \cup \underline{\beta})$ is a collection of hyperbolic right-angled hexagons, each with three non-adjacent sides that correspond to arcs of $\underline{\alpha} \cup \underline{\gamma} \cup \underline{\beta}$. The lengths of these sides determine uniquely an isometry class of right-angled hexagons, which we have just proved converge to (degenerate) right-angled hexagons in which the edges corresponding to arcs of γ become spikes in the limit. The (degenerate) right-angled hexagons glue along $\underline{\alpha}$ to form a complete hyperbolic surface Y homeomorphic to Σ with a maximal filling arc system labeled by $\underline{\alpha}$ and realized orthogeodesically on Y . That is, we have constructed a surface $Y(\underline{A}) = Y \in \mathcal{T}(\Sigma)$.

Lemma 6.8. We have an equality $\underline{A}(Y(\underline{A})) = \underline{A}$.

Proof. By construction, the length of the projection of every edge of the spine of X_t dual to an arc of $\underline{\alpha}$ was constant along the sequence (X_{t_k}) converging geometrically to $Y(\underline{A})$. The lemma follows. \square

In order to show that the inverse $Y(\cdot)$ is well-defined, we will need the following statement, which refines the relationship between the coefficients of \underline{B}_t and the lengths of its arcs.

Let $\underline{\delta} = \underline{\alpha} \cup \underline{\gamma} \cup \underline{\beta}$ denote the support of \underline{B}_t . According to [Luo07, Theorem 1.2], the lengths of the closest-point projections of the edges of the spine to dual the arcs of $\underline{\delta}$ (i.e., the coefficients of the weighted arc system) extend to an analytic local diffeomorphism $\underline{B}_\delta : \mathcal{T}(D\Sigma) \rightarrow \mathbb{R}^\delta$ whose image is a convex cone with finitely many sides.⁸ Now we show that analyticity extends to infinity.

Lemma 6.9. For each maximal filling arc system $\underline{\alpha}$ defining a cell of full dimension in $\mathcal{A}_{\text{fill}}(\Sigma, \partial\Sigma)$, there is an analytic map

$$\underline{A}_\alpha : \mathcal{T}(\Sigma) \rightarrow \mathbb{R}^\alpha$$

such that if the spine of $Y \in \mathcal{T}(\Sigma)$ has dual arc system contained in $\underline{\alpha}$, then $\underline{A}_\alpha(Y) = \underline{A}(Y)$.

Proof. The orthogeodesic length functions associated to our maximal arc system $\underline{\delta} = \underline{\alpha} \cup \underline{\gamma} \cup \underline{\beta}$ on $D\Sigma$ form an analytic parameterization of $\mathcal{T}(D\Sigma)$, which we denote by $\ell_\delta : \mathcal{T}(D\Sigma) \rightarrow \mathbb{R}_{>0}^\delta$. We have a commutative diagram of analytic embeddings

$$(9) \quad \begin{array}{ccc} \mathbb{R}_{>0}^\delta & \xrightarrow{\quad \underline{B}_\delta \circ \ell_\delta^{-1} \quad} & \mathbb{R}^\delta \\ \ell_\delta \swarrow & & \searrow \underline{B}_\delta \\ & \mathcal{T}(D\Sigma). & \end{array}$$

An explicit formula for $\underline{B}_\delta \circ \ell^{-1}$ can be recovered from [Mon09b, Section 2.4], which produces an analytic mapping $G : \mathbb{R}_{>0}^{\underline{\alpha} \cup \underline{\beta}} \rightarrow \mathbb{R}^{\underline{\alpha} \cup \underline{\beta}}$ that describes how \underline{B}_δ behaves when the arcs corresponding to $\underline{\gamma}$ have length

⁸The “projection length” associated to each arc of $\underline{\delta}$ (called the “radius coordinate” in [Luo07] and the “width” in [Mon09b]) is positive when that arc is dual to an edge of the spine of a surface $X \in \mathcal{T}(D\Sigma)$.

close to 0. More precisely, let $\pi_{\underline{\alpha} \cup \underline{\beta}} : \mathbb{R}^{\underline{\delta}} \rightarrow \mathbb{R}^{\underline{\alpha} \cup \underline{\beta}}$ be the coordinate projection. Then for $x_{\underline{\delta}} = (x_{\underline{\alpha}}, x_{\underline{\gamma}}, x_{\underline{\beta}}) \in \mathbb{R}_{>0}^{\underline{\alpha} \cup \underline{\beta}} \times \mathbb{R}_{\geq 0}^{\underline{\gamma}}$, we have

$$(10) \quad \pi_{\underline{\alpha} \cup \underline{\beta}} \circ B_{\underline{\delta}} \circ \ell_{\underline{\delta}}^{-1}(x_{\underline{\delta}}) = G(x_{\underline{\alpha}}, x_{\underline{\beta}}) + E$$

uniformly on compact subsets of $\mathbb{R}_{>0}^{\underline{\alpha} \cup \underline{\beta}} \times \mathbb{R}_{\geq 0}^{\underline{\gamma}}$, where E is a vector whose entries are all of order $O\left(\max_{\gamma \in \underline{\gamma}}\{x_{\gamma}^2\}\right)$.

Restricting to the locus of symmetric surfaces $\{X \in \mathcal{T}(D\Sigma) : \ell_{\alpha_i}(X) = \ell_{\beta_i}(X), \forall i\}$, the map G therefore induces an analytic map $F : \mathbb{R}_{>0}^{\underline{\alpha}} \rightarrow \mathbb{R}^{\underline{\alpha}}$. Again, we have an analytic parameterization $\ell_{\underline{\alpha}} : \mathcal{T}(\Sigma) \rightarrow \mathbb{R}_{>0}^{\underline{\alpha}}$ by length functions and a diagram

$$(11) \quad \begin{array}{ccc} \mathbb{R}_{>0}^{\underline{\alpha}} & \xrightarrow{F} & \mathbb{R}^{\underline{\alpha}} \\ \ell_{\underline{\alpha}} \swarrow & & \searrow F \circ \ell_{\underline{\alpha}} \\ & \mathcal{T}(\Sigma). & \end{array}$$

So take $\underline{A}_{\underline{\alpha}} = F \circ \ell_{\underline{\alpha}}$; it follows from the definitions that if the dual arc system to the spine of a surface $Y \in \mathcal{T}(\Sigma)$ is contained in $\underline{\alpha}$, then $\underline{A}_{\underline{\alpha}}(Y) = \underline{A}(Y)$. This completes the proof of the lemma. \square

A priori, $Y(\underline{A})$ depends on the subsequence X_{t_k} converging geometrically to $Y(\underline{A})$. However,

Lemma 6.10. The limit $Y(\underline{A})$ does not depend on choice of subsequence X_{t_k} , i.e., $X_t \rightarrow Y$. Moreover, $Y : |\mathcal{A}_{\text{fill}}(\Sigma, \partial\Sigma)|_{\mathbb{R}} \rightarrow \mathcal{T}(\Sigma)$ is continuous.

Proof. Throughout this proof, we let $\pi := \pi_{\underline{\alpha} \cup \underline{\beta}}$ be the coordinate projection from the proof of Lemma 6.9.

Let $s > 0$ and $X_{t,s} \in \mathcal{T}(D\Sigma)$ be the surface obtained from X_t by keeping all lengths of arcs of $\underline{\alpha} \cup \underline{\beta}$ fixed and taking $\ell_{\gamma_i}(X_{t,s}) := \ell_{\gamma_i}(X_{t+s})$, for each $\gamma_i \in \underline{\gamma}$. Note that $\ell_{\gamma_i}(X_{t+s}) = O(e^{-(s+t)/2})$. By construction of $X_{t,s}$, the lengths of arcs of $\underline{\alpha} \cup \underline{\beta}$ agree with those of X_t , so (10) gives

$$\pi(B_{\underline{\delta}}(X_t)) - \pi(B_{\underline{\delta}}(X_{t,s})) = O(e^{-(s+t)}).$$

Recall that $\pi(B_{\underline{\alpha}}(X_t)) = \pi(B_t)$ is constant for all $t > 0$, so that

$$\pi(B_{\underline{\delta}}(X_{s+t})) - \pi(B_{\underline{\delta}}(X_{t,s})) = O(e^{-(s+t)}),$$

as well. Since $B_{\underline{\delta}}$ is open, analytic, and $\{\pi(\ell_{\underline{\delta}}(X_t)) : t > 0\} \subset \mathbb{R}_{>0}^{\underline{\alpha} \cup \underline{\beta}}$ lies in a compact set (Lemmas 6.6 and 6.7), we can adjust the lengths of arcs α_i and β_i of $\underline{\alpha} \cup \underline{\beta}$ in $X_{t,s}$ by $O(e^{-(s+t)})$ to obtain X_{s+t} . Thus, for any $t_k \rightarrow \infty$, the lengths $(\ell_{t_k}(\underline{\alpha} \cup \underline{\beta}))$ form a Cauchy sequence, hence converge. Thus any two subsequential geometric limits (with basepoints away from the spikes of the subsurface associated with Σ°) coincide, which proves that $Y(\underline{A})$ is well defined.

To see that $Y(\cdot)$ is continuous, let $\underline{A}_k \rightarrow \underline{A}$; by passing to a subsequence, we may assume that \underline{A}_k are in the closure of the cell associated to a maximal filling arc system $\underline{\alpha}$. Let $\overline{\underline{A}}_k$ and $\overline{\underline{A}}$ be the mirror images (with corresponding weights) of \underline{A}_k and \underline{A} in $D\Sigma$, respectively. We build two families of approximating surfaces $X_k, X_k^k \in \mathcal{T}(D\Sigma)$ corresponding to the weighted arc systems

$$\overline{\underline{A}} + k \sum \gamma_i + \underline{A} \text{ and } \overline{\underline{A}}_k + k \sum \gamma_i + \underline{A}_k$$

on $D\Sigma$, respectively. By [Luo07, Theorem 1.2] (alternately, the proof of Lemma 6.9), each X_k^k is close to X_k in $\mathcal{T}(D\Sigma)$, hence X_k^k and X_k have the same geometric limit $Y(\underline{A}) \in \mathcal{T}(\Sigma)$, which is what we wanted to show. \square

We now have all of the pieces in place to complete the proof of Theorem 6.4.

Proof of Theorem 6.4. By Lemma 6.10, $Y(\cdot)$ is well defined and continuous, and by Lemma 6.8, $Y(\cdot)$ is a right inverse to $\underline{A}(\cdot)$; in particular, $Y(\cdot)$ is injective. For a given maximal arc system $\underline{\alpha}$, the open orthant $U_{\underline{\alpha}} = \mathbb{R}_{>0}^{\underline{\alpha}} \subset \mathbb{R}^{\underline{\alpha}}$ is identified with the interior of a top dimensional cell of $|\mathcal{A}_{\text{fill}}(\Sigma, \partial\Sigma)|_{\mathbb{R}}$. Some of the hyperplanes in $\partial\mathbb{R}_{\geq 0}^{\underline{\alpha}}$ are identified with the interior of cells associated with non-maximal filling arc systems contained in $\underline{\alpha}$; let $\overline{U}_{\underline{\alpha}}$ denote the closure of $U_{\underline{\alpha}}$ in $|\mathcal{A}_{\text{fill}}(\Sigma, \partial\Sigma)|_{\mathbb{R}}$.

Then $Y(\cdot)$ defines a continuous bijection $\overline{U_\alpha} \rightarrow \overline{Y(U_\alpha)}$, and this identification is homeomorphic, because A_α supplies an analytic inverse on $\overline{Y(U_\alpha)}$, by Lemma 6.9. Since these homeomorphisms glue along the combinatorics of $|\mathcal{A}_{\text{fill}}(\Sigma, \partial\Sigma)|_{\mathbb{R}}$, the map $Y(\cdot)$ is the desired global homeomorphic inverse to $A(\cdot)$.

Again, by Lemma 6.9, $A(\cdot)$ is analytic restricted to the relative interior of the image under $Y(\cdot)$ of each cell of $|\mathcal{A}_{\text{fill}}(\Sigma, \partial\Sigma)|_{\mathbb{R}}$, demonstrating the stratified real analytic structure. That level sets of the residue functions are mapped to one another is an exercise in unpacking the definitions. \square

7. TRANSVERSE AND SHEAR-SHAPE COCYCLES

We now define the main protagonists of this paper, the *shear-shape cocycles* on a measured lamination. In Section 7.2, we give a first definition of shear-shape cocycles in terms of the cohomology of an augmented neighborhood of λ , twisted by its local orientation (Definition 7.5). While this definition has technical merit (and exactly parallels the construction of period coordinates for quadratic differentials, a fact which we exploit in Section 10), it is impractical to use. We rectify this deficiency in Section 7.3 by giving a second formulation which parallels Bonahon's axiomatic approach to transverse cocycles (compare Definitions 7.4 and 7.12). The main result of this section, Proposition 7.14, proves that these two definitions agree.

The reader may find it helpful to consult Sections 10 or 13 while digesting these definitions so as to have a concrete model of shear-shape cocycles in mind.

7.1. Transverse cocycles. As shear-shape cocycles generalize Bonahon's transverse cocycles, we begin by recalling two equivalent definitions of transverse cocycles for geodesic laminations which we generalize in Sections 7.2 and 7.3 below.

Remark 7.1. We have chosen to present transverse cocycles in a way that anticipates our construction of shear-shape cocycles. The reader is advised that our treatment is ahistorical, and in particular omits the fascinating (and quite subtle) relationship between transverse cocycles and transverse Hölder distributions. For more on this correspondence, see [Bon97a], [Bon97b], and [Bon96].

The first definition we consider is cohomological. Let λ be a measured lamination on S ; then an orientation of λ is a continuous choice of orientation of the leaves of λ . If N is any snug neighborhood of λ , then one may take a corresponding (snug) neighborhood \widehat{N} of the orientation cover $\widehat{\lambda}$ of λ . Let ι be the covering involution of $\widehat{N} \rightarrow N$, and let $H^1(\widehat{N}, \partial\widehat{N}; \mathbb{R})^-$ denote the -1 eigenspace for the action of ι^* ,

Definition 7.2. With all notation as above, a *transverse cocycle* for λ is an element of $H^1(\widehat{N}, \partial\widehat{N}; \mathbb{R})^-$. We use $\mathcal{H}(\lambda)$ to denote the set of all transverse cocycles for λ .

With the definition above it is clear that $\mathcal{H}(\lambda)$ is a vector space, and if λ is a union of sublaminations $\lambda_1, \dots, \lambda_L$, then the space of transverse cocycles splits as

$$\mathcal{H}(\lambda) = \bigoplus_{\ell=1}^L \mathcal{H}(\lambda_\ell).$$

We record the dimension of $\mathcal{H}(\lambda)$ below.

Lemma 7.3 (Theorem 15 of [Bon97b]). The space of transverse cocycles forms a vector space of real dimension $-\chi(\lambda) + n_0(\lambda)$, where $n_0(\lambda)$ is the number of orientable components of λ .

When working with individual transverse cocycles, the above definition is rather unwieldy. Instead, it is often more useful to think of a transverse cocycle as a function on actual arcs instead of on homology classes.

Definition 7.4. Let $\lambda \in \mathcal{ML}(S)$. A *transverse cocycle* σ for λ is a function which assigns to every arc k transverse to λ a real number $\sigma(k)$ such that

- (H0) (support): If k does not intersect λ then $\sigma(k) = 0$.
- (H1) (transverse invariance): If k and k' are isotopic transverse to λ then $\sigma(k) = \sigma(k')$.
- (H2) (finite additivity): If $k = k_1 \cup k_2$ where k_i have disjoint interiors then $\sigma(k) = \sigma(k_1) + \sigma(k_2)$.

The reader familiar with train tracks will recognize that these rules resemble those governing weight systems on train tracks; see Section 9 for a continuation of this discussion.

We direct the reader to [Bon97b] or [Bon96, §3] for a proof of the equivalence of Definitions 7.2 and 7.4 (our proof of Proposition 7.14, the corresponding statement for shear-shape cocycles, can also be adapted to prove this equivalence).

7.2. Shear-shape cocycles as cohomology classes. Our first definition of a shear-shape cocycle is as a cohomology class on an appropriate augmented orientation cover, paralleling Definition 7.2. This viewpoint allows us to deduce global structural results about spaces of shear-shape cocycles (Lemma 7.8) and also reveals implicit constraints on the structure of individual shear-shape cocycles (Lemma 7.9).

Suppose that $\underline{\alpha}$ is a filling arc system for $S \setminus \lambda$. For each arc $\alpha_i \in \underline{\alpha}$, choose an arc t_i which meets α_i exactly once and is disjoint from $\lambda \cup \underline{\alpha} \setminus \{\alpha_i\}$. We call such an arc t_i a *standard transversal* to α_i . Compare Figure 9 below. An *orientation* of $\lambda \cup \underline{\alpha}$ is a continuous orientation of the leaves of λ together with a choice of orientation on each t_i such that t_i can be isotoped transverse to α_i into λ so that the orientations agree. Most pairs $\lambda \cup \underline{\alpha}$ are not orientable, but each has an *orientation double cover* $\widehat{\lambda} \cup \widehat{\underline{\alpha}}$ (the reader should have in mind the orientation cover of a quadratic differential). We note that if $\lambda \cup \underline{\alpha}$ is orientable then λ itself must be.

Consider a snug neighborhood $\mathcal{N}_\epsilon(\lambda)$ of λ on some hyperbolic surface X ; since $X \setminus \lambda$ and $X \setminus \mathcal{N}_\epsilon(\lambda)$ have the same topological type, we can identify the arc system $\underline{\alpha}$ as an arc system on $X \setminus \mathcal{N}_\epsilon(\lambda)$. In particular, taking a small neighborhood $\mathcal{N}_\epsilon(\underline{\alpha})$ of $\underline{\alpha}$ we see that there is a correspondence between complementary components of $X \setminus (\lambda \cup \underline{\alpha})$ and $X \setminus \mathcal{N}_\epsilon((\lambda \cup \underline{\alpha}))$. We will refer to any neighborhood $N_{\underline{\alpha}}$ of $\lambda \cup \underline{\alpha}$ whose complementary components have the same topological type as $X \setminus (\lambda \cup \underline{\alpha})$ as a *snug neighborhood*.

Now let $N_{\underline{\alpha}}$ be a snug neighborhood of $\lambda \cup \underline{\alpha}$; then the cover $\widehat{\lambda} \cup \widehat{\underline{\alpha}} \rightarrow \lambda \cup \underline{\alpha}$ extends to a covering $\widehat{N}_{\underline{\alpha}} \rightarrow N_{\underline{\alpha}}$ with covering involution ι . By definition of the orientation cover, each standard transversal t_i lifts to a pair of distinguished homology classes

$$t_i^{(1)}, t_i^{(2)} \in H_1(\widehat{N}_{\underline{\alpha}}, \partial \widehat{N}_{\underline{\alpha}}; \mathbb{R})$$

such that $\iota_* t_i^{(1)} = -t_i^{(2)}$.

The odd cocycles $H^1(\widehat{N}_{\underline{\alpha}}, \partial \widehat{N}_{\underline{\alpha}}; \mathbb{R})^-$ for the covering involution ι^* now provide a local cohomological model for the space of shear-shape cocycles on λ . Observe that for each i and each $\sigma \in H^1(\widehat{N}_{\underline{\alpha}}, \partial \widehat{N}_{\underline{\alpha}}; \mathbb{R})^-$, we have

$$\sigma(t_i^{(1)}) = -\iota^* \sigma(t_i^{(1)}) = -\sigma(\iota_* t_i^{(1)}) = \sigma(t_i^{(2)}).$$

Definition 7.5. Let $\lambda \in \mathcal{ML}(S)$. A *shear-shape cocycle* for λ is a pair $(\underline{\alpha}, \sigma)$ where $\underline{\alpha} = \sum \alpha_i$ is a filling arc system on $S \setminus \lambda$ and $\sigma \in H^1(\widehat{N}_{\underline{\alpha}}, \partial \widehat{N}_{\underline{\alpha}}; \mathbb{R})^-$ is such that the values $\sigma(t_i^{(j)})$ are all positive.⁹

Let $\Sigma_1 \cup \dots \cup \Sigma_m$ denote the components of $S \setminus \lambda$; then we define the *weighted arc system underlying* σ

$$\underline{A} := \sum \sigma(t_i^{(j)}) \alpha_i \in \prod_{j=1}^m |\mathcal{A}_{\text{fill}}(\Sigma_j, \partial \Sigma_j)|_{\mathbb{R}}.$$

We denote the set of all shear-shape cocycles for λ by $\mathcal{SH}(\lambda)$, the set of all shear-shape cocycles with underlying arc system $\underline{\alpha}$ by $\mathcal{SH}^\circ(\lambda; \underline{\alpha})$, and the set of all shear-shape cocycles with underlying weighted arc system \underline{A} by $\mathcal{SH}(\lambda; \underline{A})$. Often, we will leave the arc system implicit and just say that σ is a shear-shape cocycle for λ .

Remark 7.6. By Theorem 6.4, a filling weighted arc system \underline{A} is the same data as a marked hyperbolic structure on each component of $S \setminus \lambda$. In Sections 12–15 below, we prove that (so long as σ satisfies a positivity condition) these metrics glue together to give a complete hyperbolic metric on S .

Our definition of shear-shape cocycle *a priori* depends on the choice of auxiliary neighborhood $N_{\underline{\alpha}}$ of $\lambda \cup \underline{\alpha}$. However, it is not hard to see that

Lemma 7.7. The spaces of shear-shape cocycles defined by different snug neighborhoods are linearly isomorphic. Moreover, any two choices of snug neighborhoods define the same underlying weighted arc system.

Proof. Given two nested, snug neighborhoods $N'_{\underline{\alpha}} \subset N_{\underline{\alpha}}$ there is a deformation retraction of $N_{\underline{\alpha}}$ onto $N'_{\underline{\alpha}}$ (this comes from the assumption of snugness). This induces an isomorphism

$$(12) \quad H^1(\widehat{N}_{\underline{\alpha}}, \partial \widehat{N}_{\underline{\alpha}}; \mathbb{R}) \cong H^1(\widehat{N}'_{\underline{\alpha}}, \partial \widehat{N}'_{\underline{\alpha}}; \mathbb{R})$$

⁹By Poincaré–Lefschetz duality, we have a linear isomorphism $H^1(\widehat{N}_{\underline{\alpha}}, \partial \widehat{N}_{\underline{\alpha}}; \mathbb{R}) \cong H_1(\widehat{N}_{\underline{\alpha}}; \mathbb{R})$ mapping the odd cocycles for ι^* to the odd cycles for ι_* . Compare with [BD17, §§4.1 and 4.4], where a theory of (appropriately generalized) transverse (co-)cycles are applied to give shear-type coordinates for some higher rank Teichmüller spaces.

which also identifies the -1 eigenspaces of the covering involution. Therefore, we may identify the shear-shape cocycles defined by $N_{\underline{\alpha}}$ with those defined by $N'_{\underline{\alpha}}$. To see that the weights on $\underline{\alpha}$ do not depend on the choice of $N_{\underline{\alpha}}$, we note that the deformation retraction of $N_{\underline{\alpha}}$ onto $N'_{\underline{\alpha}}$ takes standard transversals to standard transversals, and hence the value of the cocycle on the transversals does not change as we change neighborhoods.

Now given any two snug neighborhoods $N_{\underline{\alpha}}$ and $N'_{\underline{\alpha}}$ of $\lambda \cup \underline{\alpha}$, one may take a common refinement $N''_{\underline{\alpha}}$ of $N_{\underline{\alpha}}$ and $N'_{\underline{\alpha}}$ and apply (12) to deduce that the spaces of shear-shape cocycles defined by $N_{\underline{\alpha}}$ and $N'_{\underline{\alpha}}$ are linearly isomorphic and define the same underlying arc system. \square

In view of this lemma, throughout the sequel we will change the neighborhood $N_{\underline{\alpha}}$ carrying σ at will.

As the orientation cover of λ naturally embeds into $\widehat{N}_{\underline{\alpha}}$, we may identify $\mathcal{H}(\lambda)$ with a subspace of $H^1(\widehat{N}_{\underline{\alpha}}, \partial \widehat{N}_{\underline{\alpha}}; \mathbb{R})$. Since any element of $\mathcal{H}(\lambda)$ evaluates to 0 on each standard transversal, we can add and subtract transverse cocycles from shear-shape cocycles without changing the underlying weighted arc system. We therefore have the following analogue of Lemma 7.3:

Lemma 7.8. Let \underline{A} be the weighted arc system underlying some shear-shape cocycle. Then $\mathcal{SH}(\lambda; \underline{A})$ is an affine space modeled on the vector space $\mathcal{H}(\lambda)$. In particular, $\dim_{\mathbb{R}}(\mathcal{SH}(\lambda; \underline{A})) = -\chi(\lambda) + n_0(\lambda)$.

Homological constraints on residues. When λ is orientable (or more generally, contains orientable components), there are homological constraints governing which weighted arc systems may underlie a shear-shape cocycle. Passing between arc systems and hyperbolic structures on complementary subsurfaces (via Theorem 6.4), these homological constraints govern when two structures can be glued together along λ .

For example, if λ is a simple closed curve then in order to glue a hyperbolic structure on $S \setminus \lambda$ along λ , the lengths of the boundary components must have equal length. Tracing through the combinatorialization by weighted arc systems, this implies that the \underline{A} -length of the boundary components must be the same. The following lemma generalizes this observation to the case when $S \setminus \lambda$ has crowned boundary (compare Lemma 13.1 below for a similar discussion using hyperbolic geometry).

Lemma 7.9. Suppose that σ is a shear-shape cocycle for λ with underlying weighted arc system \underline{A} , and let μ be an orientable component of λ . Then the sum of the (signed) residues of the boundary components incident to μ is 0.

Proof. For any component μ of λ , let $\partial(\mu)$ denote the boundary components (either closed or crowned) resulting from cutting along μ . For the purposes of this proof, let $\alpha(\mu)$ denote the sub-arc system of $\underline{\alpha}$ consisting of those arcs with endpoints on μ .

Pick an orientation on μ ; this induces an orientation on each boundary component $\mathcal{C} \in \partial(\mu)$, and hence gives the metric residue of each such \mathcal{C} a definite choice of sign. Since we are eventually going to prove that the sum of these residues is 0, it does not matter which orientation of μ we pick.

As μ is orientable, picking an orientation on μ is also equivalent to picking one of the lifts $\hat{\mu}$ of μ in the orientation cover $\widehat{\lambda} \cup \underline{\hat{\alpha}}$. Let $\widehat{\alpha(\mu)}$ denote the set of all lifts of arcs of $\alpha(\mu)$ which meet $\hat{\mu}$. Then since severing $\widehat{\alpha(\mu)}$ disconnects $\hat{\mu}$ from the rest of $\widehat{\lambda} \cup \underline{\hat{\alpha}}$, there is a relation

$$\sum_{\hat{\alpha}_i \in \widehat{\alpha(\mu)}} \varepsilon_i \hat{t}_i = 0 \text{ in } H_1(\widehat{N}_{\underline{\alpha}}, \partial \widehat{N}_{\underline{\alpha}}; \mathbb{Z})$$

where ε_i is 1 if $\hat{\alpha}_i$ is on the left-hand side of $\hat{\mu}$ and -1 on the right-hand side, and \hat{t}_i is the (relative homology class of the) oriented standard transversal corresponding to $\hat{\alpha}_i$. See Figure 8.

Therefore, for any cohomology class $\sigma \in H^1(\widehat{N}_{\underline{\alpha}}, \partial \widehat{N}_{\underline{\alpha}}; \mathbb{Z})$, and in particular any shear-shape cocycle,

$$(13) \quad \sum_{\hat{\alpha}_i \in \widehat{\alpha(\mu)}} \varepsilon_i \sigma(\hat{t}_i) = 0.$$

Now ε_i is positive when the arc is on the left-hand side of $\hat{\mu}$, equivalently (equipping $\mu \subset S$ with the corresponding orientation) when $S \setminus \lambda$ is on the left-hand side of μ . Similarly, ε_i is negative when the complementary subsurface lies to the right of μ . Unraveling the definitions and partitioning the arcs of $\alpha(\mu)$

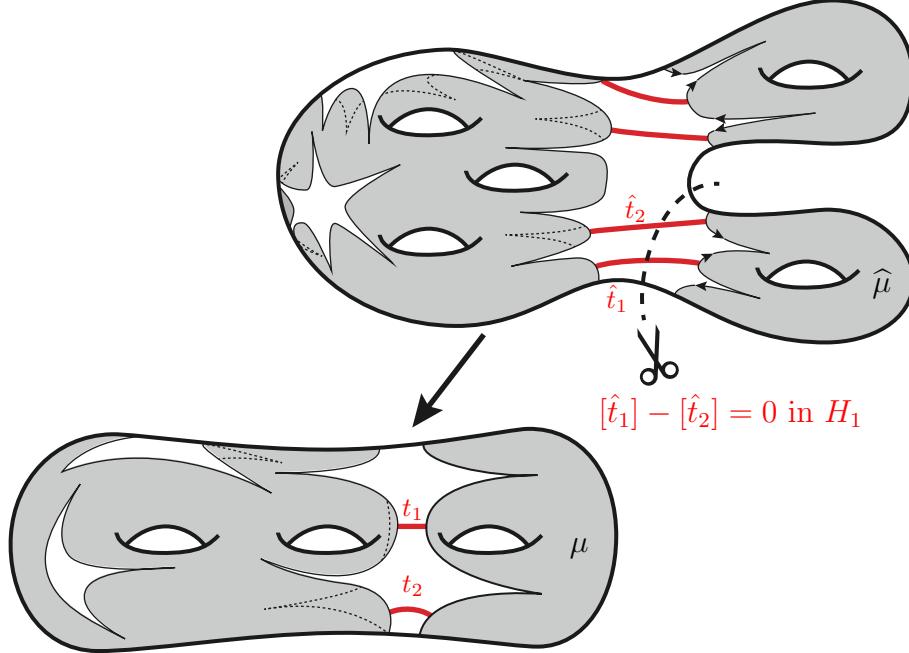


FIGURE 8. Severing ties with one of the lifts $\hat{\mu}$ of an orientable component μ of λ . This partition induces a relation in homology, hence a restriction on shear-shape cocycles. In this figure the top surface contains $\hat{\lambda}$ while the bottom contains λ ; the shaded regions are neighborhoods of these laminations.

into their incident boundary components, we see that (13) is equivalent to the statement that

$$\sum_{\mathcal{C} \in \partial(\mu)} \text{res}_{\underline{A}}(\mathcal{C}) = \sum_{\alpha_i \in \alpha(\mu)} \varepsilon_i c_i = 0,$$

which is what we wanted to prove. \square

Remark 7.10. As with transverse cocycles, one can define shear-shape cocycles for any geodesic lamination, not just those which support transverse measures. The analogue of Lemma 7.9 is more complicated in this case, as the corresponding homological relations may involve both the \hat{t}_i and other relative cycles (for example, consider when λ contains a geodesic spiraling onto a closed leaf). We have omitted such a discussion as this level of generality will not be needed for our purposes.

7.3. Shear-shape cocycles as functions on arcs. In analogy with Definition 7.4, we can also view shear-shape cocycle as functions on transverse arcs which satisfy certain properties. While this definition is more involved, it is more convenient for the calculations of Sections 13–15 and better reflects the process of “measuring” arcs by a shear-shape cocycle.

As indicated by Lemma 7.9, we must first cut out the space of all possible weighted arc systems underlying a shear-shape cocycle. Denote the complementary subsurfaces of $\lambda \in \mathcal{ML}(S)$ by $\Sigma_1, \dots, \Sigma_m$, and set

$$\mathcal{B}(S \setminus \lambda) := \left\{ \underline{A} \in \prod_{j=1}^m |\mathcal{A}_{\text{fill}}(\Sigma_j, \partial \Sigma_j)|_{\mathbb{R}} \mid \sum_{\mathcal{C} \in \partial(\mu)} \text{res}_{\underline{A}}(\mathcal{C}) = 0 \text{ for all orientable components } \mu \subset \lambda \right\}$$

where we recall that $\partial(\mu)$ denotes the set of boundary components of $S \setminus \lambda$ resulting from cutting along μ .

By Theorem 6.4, we can reinterpret $\mathcal{B}(S \setminus \lambda)$ as the set of all hyperbolic structures on $S \setminus \lambda$ so that the metric residues of the boundary components resulting from any orientable component μ of λ sum to zero. We note that when each component of λ is nonorientable, $\mathcal{B}(S \setminus \lambda)$ is just the product of the Teichmüller spaces of the complementary subsurfaces. When λ is a simple closed curve, then $\mathcal{B}(S \setminus \lambda)$ consists of those metrics on $S \setminus \lambda$ where the two boundary components have the same length.

Using this reinterpretation together with Lemma 4.4, we see that $\mathcal{B}(S \setminus \lambda)$ is topologically just a cell:

Lemma 7.11. Let $\lambda \in \mathcal{ML}(S)$ with $S \setminus \lambda = \Sigma_1 \cup \dots \cup \Sigma_m$. Then $\mathcal{B}(S \setminus \lambda) \cong \mathbb{R}^d$, where

$$d = -n_0(\lambda) + \sum_{j=1}^m \dim(\mathcal{T}(\Sigma_j))$$

where $n_0(\lambda)$ is the number of orientable components of λ .

Proof. Let $\mu_1, \dots, \mu_{n_0(\lambda)}$ denote the orientable components of λ and fix an arbitrary orientation on each. Then the lemma follows from the observation that $\mathcal{B}(S \setminus \lambda)$ is a fiber bundle over

$$\prod_{i=1}^{n_0(\lambda)} \left\{ (R_k^i) \in \mathbb{R}^{|\partial(\mu_i)|} \mid \sum_k R_k^i = 0 \right\}$$

with fibers equal to

$$\left\{ [Y, f] \in \prod_{j=1}^m \mathcal{T}(\Sigma_j) \mid \text{res}(\mathcal{C}_k) = R_k^i \text{ for each } \mathcal{C}_k \in \partial(\mu_i) \right\}.$$

By Proposition 4.5, the fibers are each homeomorphic to \mathbb{R}^d , where

$$d = \left(\sum_{j=1}^m \dim(\mathcal{T}(\Sigma_j)) \right) - \left(\sum_{i=1}^{n_0(\lambda)} |\partial(\mu_i)| \right).$$

Totaling the dimensions of base and fiber give the desired result. \square

We can now present our second definition of shear-shape cocycles.

Definition 7.12. Let $\lambda \in \mathcal{ML}(S)$. A *shear-shape cocycle* for λ is a pair (σ, \underline{A}) where \underline{A} is a weighted filling arc system

$$\underline{A} = \sum_{i=1}^n c_i \alpha_i \in \mathcal{B}(S \setminus \lambda)$$

and σ is a function which assigns to every arc k transverse to λ and disjoint from $\underline{\alpha} := \cup \alpha_i$ a real number $\sigma(k)$, satisfying the following axioms:

- (SH0) (support): If k does not intersect λ then $\sigma(k) = 0$.
- (SH1) (transverse invariance): If k and k' are isotopic through arcs transverse to λ and disjoint from $\underline{\alpha}$, then $\sigma(k) = \sigma(k')$.
- (SH2) (finite additivity): If $k = k_1 \cup k_2$ where k_i have disjoint interiors, then $\sigma(k) = \sigma(k_1) + \sigma(k_2)$.
- (SH3) (\underline{A} -compatibility): Suppose that k is isotopic rel endpoints and transverse to λ to some arc which may be written as $t_i \cup \ell$, where t_i is a standard transversal and ℓ is disjoint from $\underline{\alpha}$. Then the loop $k \cup t_i \cup \ell$ encircles a unique point p of $\lambda \cap \underline{\alpha}$, and

$$\sigma(k) = \sigma(\ell) + \varepsilon c_i$$

where ε denotes the winding number of $k \cup t_i \cup \ell$ about p (where the loop is oriented so that the edges are traversed k then t_i then ℓ). See Figure 9.

While axiom (SH3) may seem convoluted upon first inspection, its entire effect is to prescribe how the value $\sigma(k)$ evolves as an endpoint of k passes through an arc of $\underline{\alpha}$. The sign change records whether the map induced by $k = t_i \cup \ell$ from the oriented simplex into S is orientation-preserving or -reversing.

Remark 7.13. In Section 9 (Proposition 9.5 in particular), we show that there exists a choice of “smoothing” for $\underline{\alpha}$ which resolves condition (SH3) into an additivity condition. This is equivalent to prescribing that an arc k may only be dragged over a point of $\lambda \cap \underline{\alpha}$ in one direction.

The equivalence between Definitions 7.5 and 7.12 is essentially the same as the equivalence of the cohomological and axiomatic definitions of transverse cocycles [Bon96, pp. 248–9]. However, the \underline{A} -compatibility condition (axiom (SH3)) contributes new technical difficulties, and so we have included a full proof for completeness.

Proposition 7.14. The cohomological and axiomatic definitions of shear-shape cocycles agree.

Proof. Suppose first that σ is a cohomological shear-shape cocycle, that is, a cohomology class of the orientation cover $\widehat{N}_{\underline{\alpha}}$ of $N_{\underline{\alpha}}$ that is anti-invariant under the covering involution and that gives positive weight to the canonical lifts of the standard transversals of each arc of a filling arc system $\underline{\alpha}$. We begin by building from σ a function f_{σ} ; the basic idea is to restrict an arc to a neighborhood of λ , resulting in a relative homology class, and to set f_{σ} to be σ evaluated on this class.

Suppose that k is any arc transverse to λ and disjoint from $\underline{\alpha}$. Choose a small neighborhood $N_{\underline{\alpha}}$ of $\lambda \cup \underline{\alpha}$ so that k meets $\partial N_{\underline{\alpha}}$ transversely and $\partial k \cap N_{\underline{\alpha}} = \emptyset$; then $k|_{N_{\underline{\alpha}}}$ is a union of arcs with endpoints on $\partial N_{\underline{\alpha}}$. Each arc k_i of $k|_{N_{\underline{\alpha}}}$ has two distinguished, oriented lifts $k_i^{(1)}$ and $k_i^{(2)}$ to $\widehat{N}_{\underline{\alpha}}$ that cross $\widehat{\lambda}$ from right to left. As in Section 7.2, these distinguished lifts satisfy

$$(14) \quad \iota_*([k_i^{(1)}]) = -[k_i^{(2)}]$$

in $H_1(\widehat{N}_{\underline{\alpha}}, \partial \widehat{N}_{\underline{\alpha}}; \mathbb{Z})$, where ι is the covering involution of $\widehat{N}_{\underline{\alpha}} \rightarrow N_{\underline{\alpha}}$. In particular $\sigma([k_i^{(1)}]) = \sigma([k_i^{(2)}])$ since σ is anti-invariant under ι . We therefore set

$$f_{\sigma}(k) := \sigma([k])$$

where $[k]$ is the homology class of either lift of $k|_{N_{\underline{\alpha}}}$ to $\widehat{N}_{\underline{\alpha}}$.

We now prove that f_{σ} satisfies the axioms of Definition 7.12:

- (SH0) If k does not intersect λ then $k|_{N_{\underline{\alpha}}}$ is empty and $[k] = 0$, implying $f_{\sigma}(k) = 0$.
- (SH1) If k and k' are isotopic through arcs transverse to λ and disjoint from $\underline{\alpha}$ then $k|_{N_{\underline{\alpha}}}$ and $k'|_{N_{\underline{\alpha}}}$ are properly isotopic. One can lift this isotopy to the orientation cover to deduce that $[k] = [k']$ for the correct choice of lifts, so $f_{\sigma}(k) = f_{\sigma}(k')$.
- (SH2) Suppose that $k = k_1 \cup k_2$; then so long as $N_{\underline{\alpha}}$ is small enough it is clear that $k|_{N_{\underline{\alpha}}} = k_1|_{N_{\underline{\alpha}}} \cup k_2|_{N_{\underline{\alpha}}}$. Therefore, since a lift of $k|_{N_{\underline{\alpha}}}$ consists of the union of lifts of $k_1|_{N_{\underline{\alpha}}}$ and $k_2|_{N_{\underline{\alpha}}}$, we see that $[k] = [k_1] + [k_2]$ and hence the corresponding equality of f_{σ} values also holds.
- (SH3) Finally, suppose that k is isotopic (rel endpoints and transverse to λ) to $\ell \cup t_i$. Without loss of generality, we assume that the restriction of each of k, ℓ, t_i to N is a single properly embedded arc (if not, simply break the arcs into smaller pieces and apply (SH1) and (SH2) repeatedly). We also assume the restrictions are all disjoint (even at their endpoints), appealing to (SH1) as necessary.

The isotopy between k and $\ell \cup t_i$ induces a map from a disk Δ to $N_{\underline{\alpha}}$ so that $\partial \Delta \subset \partial N_{\underline{\alpha}} \cup k \cup \ell \cup t_i$. Refining $N_{\underline{\alpha}}$, isotoping the arcs, and homotoping the map as necessary, we may assume that Δ embeds into $N_{\underline{\alpha}}$, and therefore must occur in one of the configurations shown in Figure 9 below.

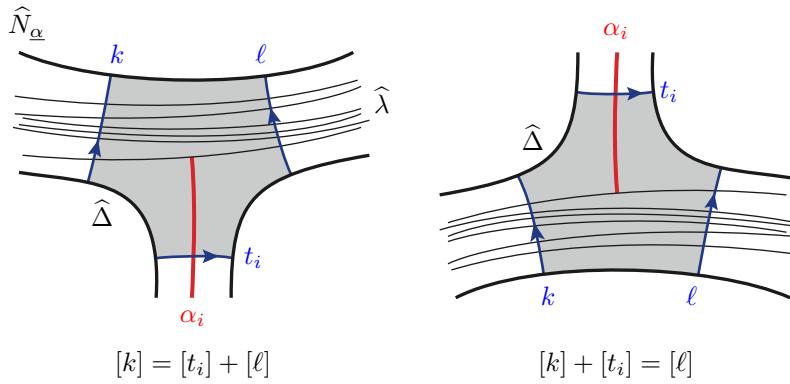


FIGURE 9. Possible configurations of the disk $\widehat{\Delta}$ and the corresponding homological relations.

Now choose one of the lifts $\widehat{\Delta} \subset \widehat{N}_{\underline{\alpha}}$ of Δ ; this choice specifies lifts of the arcs k, ℓ , and t_i and therefore (after equipping the lifts with their canonical orientations) relative homology classes $[k]$, $[\ell]$, and $[t_i]$. As these lifts together with $\partial \widehat{N}_{\underline{\alpha}}$ bound the disk $\widehat{\Delta}$, we therefore get the equality

$$[k] = [\ell] \pm [t_i]$$

where the sign is determined by the relative configuration of the arcs. Inspection of Figure 9 reveals that the sign coincides with the winding number of the loop $k \cup t_i \cup \ell$ about p .

Now suppose that (σ, \underline{A}) is an axiomatic shear-shape cocycle in the sense of Definition 7.12. Pick a snug neighborhood $N_{\underline{\alpha}}$ of $\lambda \cup \underline{\alpha}$; our task to show that the function $(k \mapsto \sigma(k))$ is indeed a cocycle (on the orientation cover, and is anti-invariant under the covering involution).

We first show that σ naturally defines a cochain on \widehat{N} relative to $\partial \widehat{N}_{\underline{\alpha}}$ which is anti-invariant by ι^* . Recall that any arc in the orientation cover comes with a canonical orientation. We may then assign to any oriented arc \hat{k} properly embedded in $\widehat{N}_{\underline{\alpha}}$ the value $\pm \sigma(k)$, where k is the image of \hat{k} under the covering projection and where the sign is positive if \hat{k} is oriented according to the canonical orientation and negative otherwise. To the (canonically oriented lifts of the) standard transversals t_i we assign the value c_i . Anti-invariance then follows by construction (compare (14)).

To see that this cochain is actually a cocycle, we show that it evaluates to 0 on every boundary. For the purposes of this argument, it will be convenient to realize $H^1(\widehat{N}_{\underline{\alpha}}, \partial \widehat{N}_{\underline{\alpha}}; \mathbb{R})$ in terms of simplicial (co)homology. The neighborhood $N_{\underline{\alpha}}$ may be triangulated as depicted in Figure 10 (compare [SB01, Figure 1]). In such a triangulation, each point of $\lambda \cap \underline{\alpha}$ and each switch of $N_{\underline{\alpha}}$ corresponds to a unique triangle, while the remaining branches each contribute a rectangle which is in turn subdivided into two triangles. This triangulation clearly lifts to an (ι -invariant) triangulation of $\widehat{N}_{\underline{\alpha}}$.

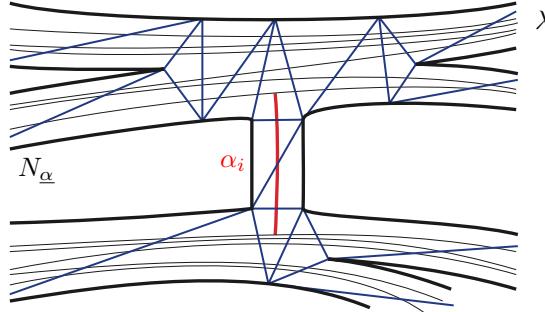


FIGURE 10. A triangulation of a (snug) neighborhood of $\lambda \cup \underline{\alpha}$. Axioms (SH0)–(SH3) imply that $\sigma(\partial\Delta) = 0$ for each triangle Δ in the triangulation, i.e., σ is a cocycle.

It therefore suffices to prove that for each oriented triangle Δ of $\widehat{N}_{\underline{\alpha}}$ we have that $\sigma(\partial\Delta) = 0$. There are three types of triangles, each of which corresponds to a different axiom of Definition 7.12.

- If Δ is (the lift of) a triangle coming from a subdivision of a branch, then one of its sides does not intersect λ and is thus assigned the value 0 by (SH0). The other two sides are isotopic rel λ , cross λ with different orientations, and are assigned the same value by (SH1). Therefore $\sigma(\partial\Delta) = 0$. Similarly, if Δ comes from a neighborhood of α , then the edges transverse to α are assigned the arc weight c_i (with opposite signs) while the other edge gets zero weight, so $\sigma(\partial\Delta) = 0$.
- Now suppose Δ is (the lift of) a triangle corresponding to a switch of $N_{\underline{\alpha}}$ with $\partial\Delta = k_1 + k_2 - k$. Then since the concatenation of k_1 and k_2 is isotopic transverse to λ to $-k$, axiom (SH2) implies

$$\sigma(k_1) + \sigma(k_2) - \sigma(k) = 0$$

and again $\sigma(\partial\Delta) = 0$.

- Finally, suppose that Δ is (the lift of) a triangle corresponding to a point of $\lambda \cap \underline{\alpha}$, so $\partial\Delta$ is some signed combination of the (canonically oriented) lifts of arcs k, ℓ and t where t is a standard transversal and k is isotopic rel endpoints and transverse to λ to $\ell \cup t$. Without loss of generality we assume that Δ is positively oriented; then depending on the configuration of k, t and ℓ we have either

$$\ell - k + t = 0 \text{ or } \ell - t - k = 0$$

(as in Figure 9). In either case, axiom (SH3) implies that $\sigma(\partial\Delta) = 0$.

We have therefore shown that $\sigma(\partial\Delta) = 0$ for every triangle of a triangulation and hence σ is indeed a 1-cocycle on $\widehat{N}_{\underline{\alpha}}$ rel boundary, finishing the proof of the lemma. \square

Measuring arcs along curves. We will also want to associate a number $\sigma(k)$ to certain arcs k that have non-empty intersection with $\underline{\alpha}$; this quantity should be invariant under suitable isotopy transverse to λ respecting the combinatorics of intersections with $\underline{\alpha}$.

So suppose $\mu \subset \lambda$ is an isolated leaf, i.e. a simple closed curve. We say that an arc k transverse to $\lambda \cup \underline{\alpha}$ and contained in an annular neighborhood of μ is *non-backtracking* if any lift \tilde{k} of k to the universal cover intersects the entire preimage $\tilde{\mu}$ of μ exactly once and \tilde{k} crosses each lift of an arc of $\underline{\alpha}$ at most once.

If k is a non-backtracking arc, then one may orient k and give μ the orientation that makes k start to the right of μ . Record the sequence of arcs β_1, \dots, β_m crossed by k , in order (note that arcs of $\underline{\alpha}$ may repeat in this sequence). Then up to isotopy, we may assume that k is a concatenation of standard transversals t_1, \dots, t_m together with a small segment k_0 disjoint from $\underline{\alpha}$ crossing μ from right to left. Compare Figure 11.

Since k is non-backtracking, the points $\beta_1 \cap \mu, \dots, \beta_m \cap \mu$ make progress around μ either in the positive direction or the negative direction. Take $\varepsilon = +1$ in the former case and $\varepsilon = -1$ in the latter, then define

$$(15) \quad \sigma(k) := \sigma(k_0) + \varepsilon \sum_{j=1}^m c_j$$

where c_j is the weight corresponding to the arc β_j . Note that the value of ε only depends on k and not on its orientation, as reversing its orientation also reverses the orientation of μ .

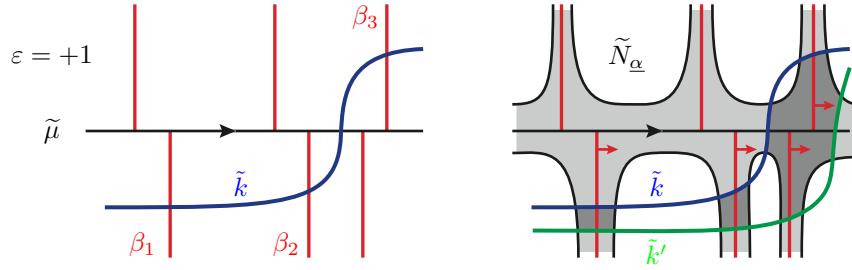


FIGURE 11. Since k makes progress around μ in the positive direction, $\varepsilon = +1$.

Lemma 7.15. Suppose that k and k' are non-backtracking arcs transverse to $\lambda \cup \underline{\alpha}$ contained in an annular neighborhood of a simple closed curve component μ of λ . If there exist lifts \tilde{k} and \tilde{k}' to \tilde{S} whose endpoints lie in the same component of $\tilde{S} \setminus (\tilde{\lambda} \cup \tilde{\alpha})$ and k is isotopic to k' transverse to λ , then $\sigma(k) = \sigma(k')$.

Proof. Fix a snug neighborhood $N_{\underline{\alpha}}$ of $\lambda \cup \underline{\alpha}$; then we need only show that $k|_{N_{\underline{\alpha}}}$ and $k'|_{N_{\underline{\alpha}}}$ define homologous cycles in the orientation cover.

We can find an isotopy $[0, 1]^2 \rightarrow \tilde{S}$ between lifts of k and k' (transverse to λ) that leaves the endpoints in the same component of $\tilde{S} \setminus \tilde{N}_{\underline{\alpha}}$. Such an isotopy then descends to S under the covering projection. The intersection of the image of each transverse arc with $N_{\underline{\alpha}}$ defines a cycle in the relative homology group, and this family of cycles is constant along the isotopy.

Since μ is orientable, an annular neighborhood of μ lifts homeomorphically to $\tilde{N}_{\underline{\alpha}}$, as do k and k' . Therefore, the isotopy between k and k' (and the homology between their restrictions) also lifts to the orientation cover $\tilde{N}_{\underline{\alpha}}$, showing that the (lifts of the) restrictions of k and k' are homologous there as well. Compare Figure 11. \square

8. THE STRUCTURE OF SHEAR-SHAPE SPACE

In this section, we investigate the global structure of the space of shear-shape cocycles. Whereas Bonahon's transverse cocycles assemble into a vector space, the space $\mathcal{SH}(\lambda)$ of all shear-shape cocycles is more complex when λ is not maximal, forming a principal $\mathcal{H}(\lambda)$ -bundle over $\mathcal{B}(S \setminus \lambda)$ (Theorem 8.1).

After understanding the structure of shear-shape space, we define an intersection form on $\mathcal{SH}(\lambda)$ (Section 8.2) and use it to specify the “positive locus” $\mathcal{SH}^+(\lambda)$ (Definition 8.4) which we show in Sections 10 through 15 serves as a global parametrization of both $\mathcal{MF}(\lambda)$ and $\mathcal{T}(S)$.

8.1. Bundle structure. Lemma 7.8 of the previous section parametrizes all shear-shape cocycles which are compatible with a given weighted arc system. In this section, we analyze how these parameter spaces piece together to get a global description of the space of all shear-shape cocycles for a fixed lamination.

Let G be a topological group. A principal G -bundle is a fiber bundle whose fibers are equipped with a transitive, continuous G -action with trivial point stabilizers together with a bundle atlas whose transition functions are continuous maps into G . We remind the reader that a principal G -bundle does not typically have a natural “zero section,” but instead, any local section of the bundle defines an identification of the fibers with G via the G -action. Moreover, any two sections define local trivializations of the bundle that differ by an element of G in each fiber.

Theorem 8.1. Let $\lambda \in \mathcal{ML}(S)$. The space $\mathcal{SH}(\lambda)$ forms a principal $\mathcal{H}(\lambda)$ -bundle over $\mathcal{B}(S \setminus \lambda)$ whose fiber over $\underline{A} \in \mathcal{B}(S \setminus \lambda)$ is $\mathcal{SH}(\lambda; \underline{A})$.

Proof. There is an obvious map from $\mathcal{SH}(\lambda)$ to $\mathcal{B}(S \setminus \lambda)$ given by remembering only the values $\sigma(t_i)$ of transversals to the arcs. For a given choice σ_0 in the fiber $\mathcal{SH}(\lambda; \underline{A})$ over \underline{A} , Lemma 7.8 identifies $\mathcal{SH}(\lambda; \underline{A})$ with $\mathcal{H}(\lambda)$ via the assignment $\sigma \mapsto \sigma - \sigma_0$.

For any filling arc system $\underline{\alpha}$ of $S \setminus \lambda$, the space $\mathcal{SH}^\circ(\lambda; \underline{\alpha})$ of shear-shape cocycles with underlying arc system $\underline{\alpha}$ is naturally identified with the open orthant

$$(16) \quad \left\{ \sigma \in H^1(\widehat{N}_{\underline{\alpha}}, \partial \widehat{N}_{\underline{\alpha}}; \mathbb{R})^- : \sigma(t_i^{(j)}) > 0 \quad \forall i, j = 1, 2 \right\},$$

where $N_{\underline{\alpha}}$ is a snug neighborhood of $\lambda \cup \underline{\alpha}$ on S .

Consider the open cell $\mathcal{B}^\circ(\underline{\alpha}) \subset \mathcal{B}(S \setminus \lambda)$ defined as all those weighted arc systems with support equal to a maximal arc system $\underline{\alpha}$. Using cohomological coordinates (16) for $\mathcal{SH}^\circ(\lambda; \underline{\alpha})$, we can find a continuous section σ of $\mathcal{SH}^\circ(\lambda; \underline{\alpha}) \rightarrow \mathcal{B}^\circ(\underline{\alpha})$. Then

$$\begin{aligned} \phi_\sigma : \mathcal{B}^\circ(\underline{\alpha}) \times \mathcal{H}(\lambda) &\rightarrow \mathcal{SH}^\circ(\lambda; \underline{\alpha}) \\ (\underline{A}, \eta) &\mapsto \sigma(\underline{A}) + \eta \end{aligned}$$

is a homeomorphism preserving fibers of the natural projections. For another choice of section σ' , we have

$$\phi_\sigma^{-1}(\phi_{\sigma'}(\underline{A}, \eta)) = (\underline{A}, \eta + \sigma'(\underline{A}) - \sigma(\underline{A})).$$

Evidently, the map $\underline{A} \mapsto \sigma'(\underline{A}) - \sigma(\underline{A}) \in \mathcal{H}(\lambda)$ is continuous.

If $N'_{\underline{\alpha}}$ is another snug neighborhood of $\lambda \cup \underline{\alpha}$, then $N_{\underline{\alpha}}$ and $N'_{\underline{\alpha}}$ share a common deformation retract. The composition of the linear isomorphisms induced on cohomology by inclusion of the deformation retract preserves the orthants defined as in (16) as well as fibers of projection to $\mathcal{B}(S \setminus \lambda)$. This proves that the principal $\mathcal{H}(\lambda)$ -structure of the bundle lying over $\mathcal{B}^\circ(\underline{\alpha})$ does not depend on the snug neighborhood whose cohomology coordinatizes $\mathcal{SH}^\circ(\lambda; \underline{\alpha})$.

To show that the principal $\mathcal{H}(\lambda)$ -bundle structures over all cells of $\mathcal{B}(S \setminus \lambda)$ glue together nicely, we find a continuous section of $\mathcal{SH}(\lambda) \rightarrow \mathcal{B}(S \setminus \lambda)$ near any given weighted arc system \underline{A} . Indeed, if $\underline{\alpha} \subset \underline{\beta}$, then inclusion $N_{\underline{\alpha}} \hookrightarrow N_{\underline{\beta}}$ of snug neighborhoods defines a map on cohomology. This map restricts to a linear isomorphism on the kernel of the evaluation map on the transversals to $\underline{\beta} \setminus \underline{\alpha}$. Thus, the closure

$$(17) \quad \mathcal{SH}(\lambda; \underline{\beta}) = \bigcup_{\substack{\underline{\beta} \supset \underline{\alpha} \\ \underline{\alpha} \text{ fills } S \setminus \lambda}} \mathcal{SH}^\circ(\lambda; \underline{\alpha})$$

of $\mathcal{SH}^\circ(\lambda; \underline{\beta})$ in $\mathcal{SH}(\lambda)$ may be realized as an orthant in $H^1(\widehat{N}_{\underline{\beta}}, \partial \widehat{N}_{\underline{\beta}}; \mathbb{R})^-$ with some open and closed faces; one of the closed faces corresponds to $\mathcal{SH}^\circ(\lambda; \underline{\alpha})$.¹⁰

Since the complex $\mathcal{A}_{\text{fill}}(S \setminus \lambda)$ is locally finite, there are only finitely many arcs β_1, \dots, β_k disjoint from $\underline{\alpha}$. Let $U \subset \mathcal{B}(S \setminus \lambda)$ be a small neighborhood of \underline{A} and σ be a continuous section of $\mathcal{SH}(\lambda; \underline{\alpha}) \rightarrow \mathcal{B}^\circ(\underline{\alpha}) \cap U$. For each i , after including $\mathcal{SH}^\circ(\lambda; \underline{\alpha})$ as a face of $\mathcal{SH}(\lambda; \underline{\alpha} \cup \beta_i)$, we may extend σ continuously on $U \cap \mathcal{B}^\circ(\underline{\alpha} \cup \beta_i)$. Continuing this process, eventually extending σ to higher dimensional cells meeting U , we end up with a continuous section $U \rightarrow \mathcal{SH}(\lambda)$, as claimed. As before, trivializations defined by two different sections differ by a continuous function $U \rightarrow \mathcal{H}(\lambda)$; this completes the proof of the theorem. \square

¹⁰When every component of $S \setminus \lambda$ is simply connected, the empty set is a filling arc system. When this is the case, $\mathcal{B}^\circ(\emptyset)$ is identified with a point, while $\mathcal{SH}(\lambda; \emptyset) = \mathcal{H}(\lambda)$.

Since every bundle over a contractible base is trivial, this implies that

Corollary 8.2. Shear-shape space $\mathcal{SH}(\lambda)$ is homeomorphic to \mathbb{R}^{6g-6} .

Proof. Let $\Sigma_1, \dots, \Sigma_m$ denote the complementary components of λ , where Σ_j has genus g_j with b_j closed boundary components and k_j crowns of types $\{c_1^j, \dots, c_{k_j}^j\}$. By Lemmas 7.11 and 4.4, we know that $\mathcal{B}(S \setminus \lambda)$ is homeomorphic to a cell of dimension

$$-n_0(\lambda) + \sum_{j=1}^m \dim(\mathcal{T}(\Sigma_j)) = -n_0(\lambda) + \sum_{j=1}^m \left(6g_j - 6 + 3b_j + \sum_{i=1}^{k_j} (c_i^j + 3) \right)$$

Lemmas 7.8 and 4.6 together imply that $\mathcal{SH}(\lambda; \underline{A})$ is an affine $\mathcal{H}(\lambda)$ -space of dimension

$$n_0(\lambda) - \chi(\lambda) = n_0(\lambda) + \frac{1}{2} \sum_{j=1}^m \sum_{i=1}^{k_j} c_i^j$$

Putting these dimension counts together via Theorem 8.1, we see that $\mathcal{SH}(\lambda)$ is homeomorphic to a cell of dimension

$$\sum_{j=1}^m \left(6g_j - 6 + 3b_j + \frac{3}{2} \sum_{i=1}^{k_j} (c_i^j + 2) \right) = \frac{3}{2\pi} \sum_{j=1}^m \text{Area}(\Sigma_j) = \frac{3}{2\pi} \text{Area}(S) = 6g - 6,$$

where the first equality follows from (4). \square

8.2. Intersection forms and positivity. Now that we have a global description of shear-shape space, we restrict our attention to a certain positive locus $\mathcal{SH}^+(\lambda)$ inside of $\mathcal{SH}(\lambda)$. The main result of this section is Proposition 8.5, in which we identify $\mathcal{SH}^+(\lambda)$ as an affine cone bundle over $\mathcal{B}(S \setminus \lambda)$.

Positive transverse cocycles. We begin by recalling the definition of positivity for transverse cocycles, as developed in [Bon96, §6]. Fixing some $\lambda \in \mathcal{ML}(S)$, we recall that a transverse cocycle for λ may be identified with a relative cohomology class of the orientation cover \hat{N} of a snug neighborhood N of λ (Definition 7.2). The intersection pairing of \hat{N} therefore induces a anti-symmetric bilinear pairing

$$\omega_{\mathcal{H}} : \mathcal{H}(\lambda) \times \mathcal{H}(\lambda) \rightarrow \mathbb{R}$$

called the *Thurston intersection/symplectic form*. This form is nondegenerate when λ is maximal, and more generally, when λ cuts S into polygons each with an odd number of sides [PH92, §3.2].

Each transverse measure for λ is in particular a transverse cocycle. Using the intersection form one can therefore define a positive cone $\mathcal{H}^+(\lambda)$ inside of $\mathcal{H}(\lambda)$ with respect to the (non-atomic) measures supported on λ . Write

$$\lambda = \lambda_1 \cup \dots \cup \lambda_L \cup \gamma_1 \cup \dots \cup \gamma_M$$

where the γ_m are all weighted simple closed curves and the λ_ℓ are minimal measured sub-laminations whose supports are not simple closed curves. Then set

$$(18) \quad \mathcal{H}^+(\lambda) := \{ \rho \in \mathcal{H}(\lambda) : \omega_{\mathcal{H}}(\rho, \mu) > 0 \text{ for all } \mu \in \bigcup_{\ell=1}^L \Delta(\lambda_\ell) \},$$

where $\Delta(\lambda_\ell)$ denotes the collection of measures supported on λ_ℓ .

The reason for this involved definition is that the Thurston form is identically 0 exactly when the underlying lamination is a multicurve. Therefore, if the support of λ contains a simple closed curve γ , the pairing of γ with every transverse cocycle supported on λ is 0.¹¹

On the other hand, so long as λ is not a multicurve then the Thurston form is not identically 0. In fact, the cone $\mathcal{H}^+(\lambda)$ splits as a product

$$\mathcal{H}^+(\lambda) = \bigoplus_{\ell=1}^L \mathcal{H}^+(\lambda_\ell) \oplus \bigoplus_{m=1}^M \mathcal{H}(\gamma_m).$$

¹¹This is because the components of the orientation cover are all annuli, whose first (co)homologies all have rank 1. For non-curve laminations, the homology has higher rank and so can support a nonzero intersection form.

As λ supports at most $3g - 3$ (projective classes of) ergodic measures, each $\mathcal{H}^+(\lambda_\ell)$ is a cone with a side for each (projective class of) ergodic measure supported on λ_ℓ .

When λ is a multicurve, then there are no λ_ℓ 's and so the condition of (18) is empty. As such, in this case we have that the space of positive transverse cocycles is the entire twist space:

$$\mathcal{H}^+(\gamma_1 \cup \dots \cup \gamma_M) = \mathcal{H}(\gamma_1 \cup \dots \cup \gamma_M) = \bigoplus_{m=1}^M \mathcal{H}(\gamma_m) \cong \mathbb{R}^M.$$

Therefore, we see that no matter whether γ is a multicurve or not, the space $\mathcal{H}^+(\lambda)$ is a convex cone of full dimension (where we expand our definition of “cone” to include the entire vector space).

Positive shear-shape cocycles. We now repeat the above discussion for shear-shape cocycles. By Definition 7.5, any shear-shape cocycle $(\sigma, \underline{\alpha})$ may be identified with a relative cohomology class of the orientation cover $\widehat{N}_{\underline{\alpha}}$ of a neighborhood $N_{\underline{\alpha}}$ of $\lambda \cup \underline{\alpha}$. As above, the intersection pairing of \widehat{N} then defines a pairing between any two shear-shape cocycles with underlying arc system contained inside of $\underline{\alpha}$. However, if the underlying arc systems of $\sigma, \rho \in \mathcal{SH}(\lambda)$ are not nested then there is no obvious way to pair the two cocycles.

While it does not make sense to pair two arbitrary shear-shape cocycles, we can always pair shear-shape cocycles with transverse cocycles. Recall from (the discussion before) Lemma 7.8 that $\mathcal{H}(\lambda)$ naturally embeds as a subspace of the cohomology of the neighborhood $\widehat{N}_{\underline{\alpha}}$ defining a shear-shape cocycle and may be identified with the kernel of the evaluation map on transversals to $\underline{\alpha}$. Therefore, the intersection pairing on $\widehat{N}_{\underline{\alpha}}$ gives rise to a function

$$\omega_{\mathcal{SH}} : \mathcal{SH}(\lambda) \times \mathcal{H}(\lambda) \rightarrow \mathbb{R}$$

which we also refer to as the *Thurston intersection form*. Throughout the paper, we will differentiate between the different intersection forms by indicating their domains in subscript.

We record some of the relevant properties of $\omega_{\mathcal{SH}}$ below.

Lemma 8.3. The Thurston intersection form $\omega_{\mathcal{SH}}$ is a $\text{Mod}(S)[\lambda]$ -invariant continuous pairing which is homogeneous in the first factor and linear in the second. Moreover, for any $\underline{A} \in \mathcal{B}(S \setminus \lambda)$ and $\rho \in \mathcal{H}(\lambda)$, the function

$$\omega_{\mathcal{SH}}(\cdot, \rho) : \mathcal{SH}(\lambda; \underline{A}) \rightarrow \mathbb{R}$$

is an affine homomorphism inducing $\omega_{\mathcal{H}}(\cdot, \rho)$ on the underlying vector space $\mathcal{H}(\lambda)$.

Proof. We begin by showing that the form is actually well-defined. Suppose first that $\underline{\alpha}$ is maximal; then since the (homological) intersection form is natural with respect to deformation retracts, and any two snug neighborhoods of $\lambda \cup \underline{\alpha}$ share a common deformation retract, we see that the form does not depend on the choice of neighborhood.

Now suppose that $\underline{\beta}$ is a filling arc system that is a subsystem of two different maximal arc systems $\underline{\alpha}_1$ and $\underline{\alpha}_2$. Then one can take a snug neighborhood $N_{\underline{\beta}}$ of $\lambda \cup \underline{\beta}$ which includes into neighborhoods N_i of $\lambda \cup \underline{\alpha}_i$ for $i = 1, 2$. Now since the (homological) intersection form is also natural with respect to inclusions, we see that the Thurston form must be as well. Therefore, for any $\sigma \in \mathcal{SH}(\lambda; \underline{\beta})$ and $\rho \in \mathcal{H}(\lambda)$ it does not matter if we compute $\omega_{\mathcal{SH}}(\sigma, \rho)$ in $N_{\underline{\beta}}$, N_1 , or N_2 .

Now that we have established that $\omega_{\mathcal{SH}}$ is well-defined, the other properties follow readily from properties of the (homological) intersection form. Since the homological intersection pairing is linear in each coordinate, we get that $\omega_{\mathcal{SH}}$ is in particular linear in the second coordinate. Similarly, for any $\underline{A} \in \mathcal{B}(S \setminus \lambda)$ and any two $\sigma_1, \sigma_2 \in \mathcal{SH}(\lambda; \underline{A})$ we know that $\sigma_1 - \sigma_2$ is a transverse cocycle, and again by linearity of the homological intersection form we get that

$$\omega_{\mathcal{SH}}(\sigma_1, \rho) - \omega_{\mathcal{SH}}(\sigma_2, \rho) = \omega_{\mathcal{H}}(\sigma_1 - \sigma_2, \rho)$$

for all $\rho \in \mathcal{H}(\lambda)$. Thus $\omega_{\mathcal{SH}}$ is affine on each $\mathcal{SH}(\lambda; \underline{A})$.

Finally, to see that the map $\omega_{\mathcal{SH}}(\cdot, \rho)$ is continuous for a fixed ρ , we recall that for any maximal arc system $\underline{\alpha}$, the space $\mathcal{SH}^o(\lambda; \underline{\alpha})$ of shear-shape cocycles with underlying arc system $\underline{\alpha}$ may be realized as an open octant in cohomological coordinates (16), and this parametrization extends to its closure $\mathcal{SH}(\lambda; \underline{\alpha})$.

Since the intersection pairing on cohomology is continuous, we see that for each maximal arc system $\underline{\alpha}$ the function $\omega_{\mathcal{SH}}(\cdot, \rho)$ is continuous on $\mathcal{SH}(\lambda; \underline{\alpha})$. But now since we have checked that the value of $\omega_{\mathcal{SH}}(\cdot, \rho)$ does not actually depend on the neighborhood, it agrees on the overlaps of closures $\mathcal{SH}(\lambda; \underline{\alpha})$ for maximal $\underline{\alpha}$.

Therefore, since the cell structure of $\mathcal{B}(S \setminus \lambda)$ is locally finite we may glue together the functions $\omega_{\mathcal{SH}}(\cdot, \rho)$ (which are continuous on each $\mathcal{SH}(\lambda; \underline{\alpha})$) to get a globally continuous function on $\mathcal{SH}(\lambda)$. \square

With this intersection form in hand, we may now define a positive locus with respect to the set of measures supported on λ .

Definition 8.4. The space of *positive shear-shape cocycles* $\mathcal{SH}^+(\lambda)$ is the set

$$\mathcal{SH}^+(\lambda) = \{\sigma \in \mathcal{SH}(\lambda) : \omega_{\mathcal{SH}}(\sigma, \mu) > 0 \text{ for all } \mu \in \Delta(\lambda)\}.$$

Observe the difference between the definition above and the one appearing in (18): any positive shear-shape cocycle must also pair positively with all simple closed curves γ_m appearing in the support of λ . The essential difference between the two cases is that additional branches of $\tau_{\underline{\alpha}}$ coming from the underlying arc system allows a shear-shape cocycle to meet each γ_m without being completely supported on γ_m . Indeed, one can check that the contribution to the Thurston form coming from the intersection of $\underline{\alpha}$ with a simple closed curve component of λ is always positive (compare (20)). In particular, the positivity condition is automatically fulfilled for any measure supported on a curve component of λ .

On each cohomological chart (16) or (17) it is clear that $\mathcal{SH}^+(\lambda)$ is an open cone cut out by finitely many linear inequalities (one for each ergodic measure supported on λ , plus positivity of arcs weights). However, this does not yield a global description of $\mathcal{SH}^+(\lambda)$. In order to get one, we must show that the linear subspaces cut out by the positivity conditions intersect the $\mathcal{H}(\lambda)$ fibers transversely.

Proposition 8.5. The space $\mathcal{SH}^+(\lambda)$ is an affine cone bundle over $\mathcal{B}(S \setminus \lambda)$ with fibers isomorphic to $\mathcal{H}^+(\lambda)$.

By an affine cone bundle, we mean that there is a (non-unique) section $\sigma_0 : \mathcal{B}(S \setminus \lambda) \rightarrow \mathcal{SH}(\lambda)$ such that

$$\mathcal{SH}^+(\lambda) \cap \mathcal{SH}(\lambda; \underline{A}) = \sigma_0(\underline{A}) + \mathcal{H}^+(\lambda)$$

for every $\underline{A} \in \mathcal{B}(S \setminus \lambda)$. Moreover, any two such sections differ by a continuous map $\mathcal{B}(S \setminus \lambda) \rightarrow \mathcal{H}(\lambda)$.

Proof. Choose mutually singular ergodic measures $\mu_1, \dots, \mu_N, \gamma_1, \dots, \gamma_M$ on λ that span $\Delta(\lambda)$, where the supports of the μ_n are non-curve laminations and the γ_m are all simple closed curves. Pick an arbitrary $\sigma \in \mathcal{SH}(\lambda; \underline{A})$, and define

$$C(\sigma) := \{\rho \in \mathcal{H}(\lambda) \mid \omega_{\mathcal{H}}(\rho, \mu_n) > -\omega_{\mathcal{SH}}(\sigma, \mu_n) \text{ for all } n = 1, \dots, N\}.$$

By linearity of $\omega_{\mathcal{H}}$ on $\mathcal{H}(\lambda)$, together with the fact that the pairing $\omega_{\mathcal{H}}(\cdot, \mu_n)$ is not identically 0 since the support of μ_n is not a simple closed curve, this is an intersection of N affine half-spaces which do not depend on our choice of ergodic measures μ_i in their projective classes. Again by linearity, we see that this is just a translate of $\mathcal{H}^+(\lambda)$ and hence is a cone of full dimension.

Now since $\omega_{\mathcal{SH}}(\cdot, \mu_j)$ is an affine map on $\mathcal{SH}(\lambda; \underline{A})$ for each j , we see that

$$\sigma + C(\sigma) = \{\eta \in \mathcal{SH}(\lambda; \underline{A}) \mid \omega_{\mathcal{SH}}(\eta, \mu_n) > 0 \text{ for all } n = 1, \dots, N\} = \mathcal{SH}^+(\lambda) \cap \mathcal{SH}(\lambda; \underline{A})$$

is an affine cone of full dimension (where the last equality holds because the positive discussion is automatically fulfilled for each γ_m). It is a further consequence of affinity that this identification does not depend on the choice of σ . The bundle structure then follows from continuity of $\omega_{\mathcal{SH}}$. \square

9. TRAIN TRACK COORDINATES FOR SHEAR-SHAPE SPACE

In this section, we introduced train track charts for shear-shape cocycles. In Section 9.1 we recall Bonahon's realization of transverse cocycles to a lamination in the weight space of a train track that snugly carries it. In Section 9.2, we reinterpret the cohomological coordinate charts (16) for $\mathcal{SH}^+(\lambda; \underline{\alpha})$ by "smoothing" $\lambda \cup \underline{\alpha}$ onto a train track $\tau_{\underline{\alpha}}$ (Construction 9.3) and realizing $\mathcal{SH}^+(\lambda; \underline{\alpha})$ as an orthant in the weight space of $\tau_{\underline{\alpha}}$ (Proposition 9.5). This construction also has the added benefit of converting axiom (SH3) of Definition 7.12 into a simpler additivity condition; this is convenient for computations and provides an explicit formula (20) for the Thurston intersection pairing. We rely on this formula in Section 10.2 to show that foliations transverse to λ define positive shear-shape cocycles (Proposition 10.12).

Later, in Section 9.3, we explain how the PIL structure of $\mathcal{SH}(\lambda)$ is manifest in train track coordinates and provides a canonical measure in the class of Lebesgue. When λ is maximal, this measure is a constant multiple of the symplectic volume element induced by $\omega_{\mathcal{H}}$. Finally, in Section 9.4 we consider how train track charts facilitate an interpretation of $\mathcal{SH}(\lambda)$ as organizing the fragments of the cotangent space to \mathcal{ML} at λ .

Remark 9.1. We advise the reader that two different types of train tracks appear below: those which carry transverse cocycles for λ and give coordinates on the fiber $\mathcal{SH}(\lambda; \underline{A})$, and those which carry shear-shape cocycles and give coordinates on the total space $\mathcal{SH}(\lambda)$.

9.1. Train track coordinates for transverse cocycles. We begin by recalling how transverse cocycles can be parametrized by weight systems on (snug) train tracks. The advantage of these coordinates is that they determine the cocycle with only finitely many values (a main benefit of the cohomological Definition 7.2), but do so using unoriented arcs on the surface, not the orientation cover (a main benefit of the axiomatic Definition 7.4).

Let τ be a train track snugly carrying a geodesic lamination λ and σ a transverse cocycle, thought of as a function on transverse arcs. For each branch b of τ , pick a tie t_b . Then one can assign to b the weight $\sigma(t_b)$; by Axiom (H1) this value does not depend on the choice of tie, and by Axiom (H2) these weights necessarily satisfy the switch conditions. Therefore, any transverse cocycle can be represented by a weight system on τ , and in fact this map is an isomorphism.

Proposition 9.2 (Theorem 11 of [Bon97b]). Let τ be a train track snugly carrying a geodesic lamination λ . Then the map $\sigma \mapsto \{\sigma(t_b)\}_{b \in b(\tau)}$ is a linear isomorphism between $\mathcal{H}(\lambda)$ and $W(\tau)$, the space of all (real) weights on τ satisfying the switch conditions.

On a given train track snugly carrying λ , the Thurston intersection form $\omega_{\mathcal{H}}$ is easily computable in terms of the weight systems. To wit, if $\sigma, \rho \in \mathcal{H}(\lambda)$ then their intersection is equal to

$$(19) \quad \omega_{\mathcal{H}}(\sigma, \rho) = \frac{1}{2} \sum_s \begin{vmatrix} \sigma(r_s) & \sigma(\ell_s) \\ \rho(r_s) & \rho(\ell_s) \end{vmatrix}$$

where the sum is over all switches s of τ and r_s and ℓ_s are the half-branches which leave s from the right and the left, respectively. Compare [PH92, §3.2].

9.2. Train track coordinates for shear-shape cocycles. In order to imitate the above construction for shear-shape cocycles, we first must explain how to build a train track from λ and a filling arc system $\underline{\alpha}$ on its complement.

Suppose that τ carries λ snugly; then the complementary components of $\tau \cup \underline{\alpha}$ correspond to those of $\lambda \cup \underline{\alpha}$. A *smoothing* of $\tau \cup \underline{\alpha}$ is a train track $\tau_{\underline{\alpha}}$ which is obtained by choosing tangential data at each of the points of $\tau \cap \underline{\alpha}$ and isotoping each arc of $\underline{\alpha}$ to meet τ along the prescribed direction. Each component of $S \setminus \tau$ inherits an orientation from S , which in turn gives an orientation to the boundary (of the metric completion) of each subsurface. A smoothing $\tau_{\underline{\alpha}}$ is *standard* if for each switch of $\tau_{\underline{\alpha}}$ with an incoming half branch corresponding to an arc $\alpha_i \in \underline{\alpha}$, the incoming tangent vector to α_i is pointing in the positive direction with respect to the boundary orientation of the component of $S \setminus \tau$ containing α_i ; see Figure 12.

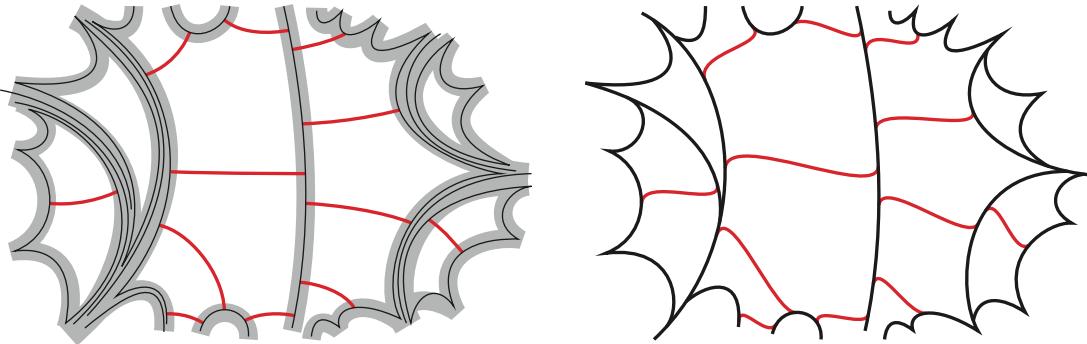


FIGURE 12. Left: a geometric train track neighborhood of $\tilde{\lambda}$ together with an arc system $\tilde{\alpha}$. Right: The (preimage of the) standard smoothing $\tau_{\underline{\alpha}}$.

Recall (Construction 5.6) that a geometric train track τ constructed from a hyperbolic structure $X \in \mathcal{T}(S)$, $\lambda \in \mathcal{ML}(S)$, and $\epsilon > 0$ is obtained as the leaf space of the orthogeodesic foliation restricted to an ϵ -neighborhood of λ in X (for small enough values of ϵ).

Construction 9.3 (Geometric standard smoothings). Let $\lambda \in \mathcal{ML}(S)$ and X be a hyperbolic metric on S . Let $\underline{\alpha}$ be a filling arc system in $S \setminus \lambda$, realized orthogeodesically on X . For small enough $\epsilon > 0$, $\underline{\alpha} \cap \mathcal{N}_\epsilon(\lambda)$ lies in a finite collection of leaves of $\mathcal{O}_\lambda(X)$ and so each end of each arc of $\underline{\alpha}$ defines a point in the quotient $\tau = \mathcal{N}_\epsilon(\lambda) / \sim$, where \sim is the equivalence relation induced by collapsing the leaves of $\mathcal{O}_\lambda(X)|_{\mathcal{N}_\epsilon(\lambda)}$.

The geometric standard smoothing $\tau_{\underline{\alpha}}$ is then obtained by attaching $\underline{\alpha}$ onto the geometric train track τ at these points and smoothing in the standard way.

Since $\underline{\alpha}$ is filling, the components of $X \setminus (\lambda \cup \underline{\alpha})$ are topological disks. In a geometric standard smoothing $\tau_{\underline{\alpha}}$, each complementary disk incident to an arc α of $\underline{\alpha}$ has at least one spike corresponding to an end of that α . Since no arc of $\underline{\alpha}$ joins asymptotic geodesics of λ , the complementary polygons all have at least 3 spikes and so we see that $\tau_{\underline{\alpha}}$ is indeed a train track.

Remark 9.4. A geometric standard smoothing keeps track of the intersection pattern of λ with $\underline{\alpha}$ on “either side” of τ , and the endpoints of $\underline{\alpha}$ on a geometric train track $\tau_\epsilon \subset X$ constructed from λ by a parameter $\epsilon > 0$ as in Construction 9.3 are stable as $\epsilon \rightarrow 0$.

A standard smoothing $\tau_{\underline{\alpha}}$ is reminiscent of the construction of completing λ to a maximal lamination λ' by “spinning” the arcs of $\underline{\alpha}$ around the boundary geodesics of complementary subsurfaces to λ in the positive direction to obtain spiraling isolated leaves of λ' in bijection with the arcs of $\underline{\alpha}$. In Proposition 9.5 below, we observe that by smoothing $\underline{\alpha}$ onto τ in a standard way, axiom (SH3) allows us to assign weights to the branches of $\tau_{\underline{\alpha}}$ in such a way that the switch conditions are satisfied. Thus, for a shear-shape cocycle carried by $\tau_{\underline{\alpha}}$, the weights deposited on the branches $\underline{\alpha} \subset \tau_{\underline{\alpha}}$ encode “shape” data, rather than “shear” data. As such, we do not think of a standard smoothing as corresponding to the completion of λ to a maximal lamination λ' .

Proposition 9.5. Every shear-shape cocycle $(\sigma, \underline{A}) \in \mathcal{SH}(\lambda)$ may be represented by a weight system $w_{\underline{\alpha}}(\sigma)$ on a standard smoothing $\tau_{\underline{\alpha}}$ that also carries λ . Moreover, the map $\sigma \mapsto w_{\underline{\alpha}}(\sigma)$ extends to a linear isomorphism

$$H^1(\widehat{N}_{\underline{\alpha}}, \partial \widehat{N}_{\underline{\alpha}}; \mathbb{R})^- \cong W(\tau_{\underline{\alpha}})$$

where $N_{\underline{\alpha}}$ is a neighborhood of $\lambda \cup \underline{\alpha}$, $\widehat{N}_{\underline{\alpha}}$ is its orientation cover, and $H^1(\widehat{N}_{\underline{\alpha}}, \partial \widehat{N}_{\underline{\alpha}}; \mathbb{R})^-$ is the -1 eigenspace for the covering involution ι^* .

In particular, this isomorphism realizes $\mathcal{SH}(\lambda; \underline{\alpha})$ and $\mathcal{SH}^+(\lambda, \underline{\alpha})$ as convex cones (with some open and some closed faces) inside of $W(\tau_{\underline{\alpha}})$.

Proof. Let $\tau_{\underline{\alpha}}$ be a standard smoothing of $\tau \cup \underline{\alpha}$ and for each branch b of $\tau_{\underline{\alpha}}$, let t_b denote a tie transverse to b . Evaluating a shear-shape cocycle σ on t_b yields an assignment of weights

$$w_{\underline{\alpha}}(\sigma) : b \rightarrow \sigma(t_b).$$

By axiom (SH1) of Definition 7.12, this weight system does not depend on the choice of tie.

To check that $w_{\underline{\alpha}}(\sigma)$ satisfies the switch conditions, we observe that there are two types of switches of $\tau_{\underline{\alpha}}$: those that come from switches of τ and those that come from smoothings of points of $\lambda \cap \underline{\alpha}$. Axiom (SH2) implies that the switch condition holds at each of the former, while axiom (SH3) together with our choice of smoothing ensures that $w_{\underline{\alpha}}(\sigma)$ satisfies the switch conditions at each of the latter. Compare Figure 13.

We note that this discussion does not rely on the positivity of σ on standard transversals, and so can be repeated to realize an arbitrary element of $H^1(\widehat{N}_{\underline{\alpha}}, \partial \widehat{N}_{\underline{\alpha}}; \mathbb{R})^-$ as a weight system on $\tau_{\underline{\alpha}}$. \square

Let $\underline{A} = \sum c_i \alpha_i$; then on any smoothing $\tau_{\underline{\alpha}}$ the identification of Proposition 9.5 restricts to an isomorphism

$$\mathcal{SH}(\lambda; \underline{A}) \cong \{w \in W(\tau_{\underline{\alpha}}) : w(b_i) = c_i\}$$

where b_i is the branch of $\tau_{\underline{\alpha}}$ corresponding to α_i . Indeed, these coordinates together with the parametrization of transverse cocycles by weight systems on $\tau \prec \tau_{\underline{\alpha}}$ (Proposition 9.2) give another proof that the difference of any two shear-shape cocycles compatible with a given $\underline{A} \in \mathcal{B}(S \setminus \lambda)$ is a transverse cocycle (Lemma 7.8).

Remark 9.6. The metric residue condition (Lemma 7.9) is still visible in train track coordinates, though it is somewhat obscured. Indeed, suppose that λ contains an orientable component carried on a component ζ of the geometric train track τ ; fix an arbitrary orientation of ζ .

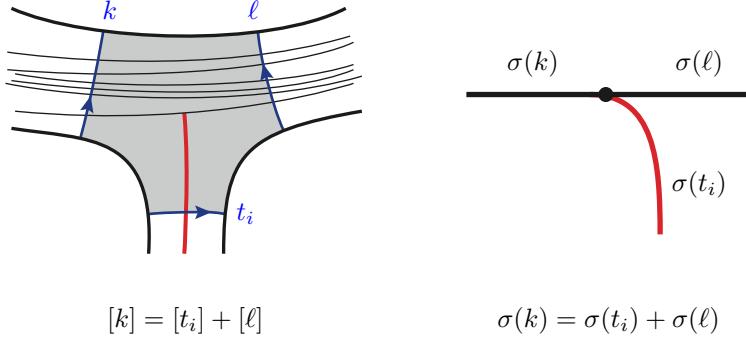


FIGURE 13. A standard smoothing of a geometric train track. The equation in homology encoded by axiom (SH3) becomes an additivity condition on the train track.

Take a geometric standard smoothing $\tau_{\underline{\alpha}}$ of $\tau \cup \underline{\alpha}$. Reversing the tangential information as necessary, we can then construct a (non-standard) smoothing of $\tau \cup \underline{\alpha}$ so that every arc of $\underline{\alpha}$ is a small branch entering ζ according to the orientation. Moreover, by reversing the sign of the weight on each arc which has its smoothing data modified, this non-standard smoothing still carries shear-shape cocycles as a weight systems. But then by conservation of mass the total sum of the weights on the branches entering ζ must be 0.

Hence in this setting the metric residue condition manifests as a condition embedded in the recurrence structure of smoothings.

The extended intersection form on $\mathcal{SH}(\lambda)$ also has a nice formula in terms of train tracks. Let τ be a (trivalent) train track snugly carrying λ and let $\tau_{\underline{\alpha}}$ be a standard smoothing of $\tau \cup \underline{\alpha}$; then for $\sigma \in \mathcal{SH}(\lambda)$ and $\rho \in \mathcal{H}(\lambda)$, we have

$$(20) \quad \omega_{\mathcal{SH}}(\sigma, \rho) = \frac{1}{2} \sum_s \begin{vmatrix} \sigma(r_s) & \sigma(\ell_s) \\ \rho(r_s) & \rho(\ell_s) \end{vmatrix}$$

where the sum is over all switches s of $\tau_{\underline{\alpha}}$ and r_s/ℓ_s are the right/left small half-branches. The proof of this formula is the same as that of (19) and is therefore omitted; the only thing to note in this case is that the value does not change if one completes $\underline{\alpha}$ by adding in arcs of zero weight.

9.3. Piecewise-integral-linear structure. A piecewise linear manifold is said to be *piecewise-integral-linear* or *PIL* with respect to a choice of charts if the transition functions are invertible piecewise-linear maps with integral coefficients. The track charts that we have constructed from standard smoothings in this section endow each cell $\mathcal{SH}(\lambda; \underline{\alpha})$ with a PIL structure which clearly extends over all of $\mathcal{SH}(\lambda)$ (compare [PH92, §3.1]).

The points of the integer lattice in $W(\tau_{\underline{\alpha}})$ are invariant under coordinate transformation, thus the *integer points* $\mathcal{SH}_{\mathbb{Z}}(\lambda) \subset \mathcal{SH}(\lambda)$ are well defined.

The PIL structure defined by train track charts gives a canonical measure $\mu_{\mathcal{SH}}$ in the class of the $(6g-6)$ -dimensional Lebesgue measure on $\mathcal{SH}(\lambda)$. Namely, if $B \subset \mathcal{SH}(\lambda)$ is a Borel set, then

$$(21) \quad \mu_{\mathcal{SH}}(B) := \lim_{R \rightarrow \infty} \frac{\#R \cdot B \cap \mathcal{SH}_{\mathbb{Z}}}{R^{6g-6}}.$$

Since the symplectic intersection form $\omega_{\mathcal{SH}}$ is constant (19) in a train track chart, the volume element defined by the $(3g-3)$ -fold wedge product $\wedge \omega_{\mathcal{SH}}$ is a constant multiple of $\mu_{\mathcal{SH}}$ on each chart.

We note that $\mathcal{B}(S \setminus \lambda)$ is cut out of $|\mathcal{A}_{\text{fill}}(S \setminus \lambda)|$ by linear equations with integer coefficients, as is each cell of $|\mathcal{A}_{\text{fill}}(S \setminus \lambda)|$. Therefore, the integer lattice $\mathcal{SH}_{\mathbb{Z}}(\lambda)$ restricts to a integer lattice in the bundle $\mathcal{SH}(\lambda; \underline{\alpha})$ over every cell $\mathcal{B}(\underline{\alpha})$. Thus we obtain a natural volume element on the bundle over the k -skeleton of $\mathcal{B}(S \setminus \lambda)$, whenever it is not empty.

9.4. Duality in train track coordinates. We now take a moment to discuss shear-shape coordinates from the point of view of train track weight spaces; this discussion is motivated by that in [Thu98], and is meant to clarify how shear-shape cocycles fit into the broader theory of train tracks.

We begin by recalling the analogy between shear coordinates for Teichmüller space and the “horospherical coordinates” for hyperbolic space. As observed by Thurston [Thu98, p. 42], projecting the Lorentz model

$$\mathbb{H}^n = \{x_1^2 + \dots + x_n^2 - x_{n+1}^2 = -1 \mid x_{n+1} > 0\}$$

to $\langle x_1, \dots, x_n \rangle$ along a family of parallel light rays gives a parametrization for \mathbb{H}^n in terms of a half-space. In these coordinates, horospheres based at the boundary point $\xi \in \partial_\infty \mathbb{H}^n$ corresponding to the choice of light ray are mapped to affine hyperplanes and geodesics from ξ are mapped to rays from the origin. ^[12]

When λ is maximal and uniquely ergodic, Bonahon and Thurston’s shear coordinates similarly realize $\mathcal{T}(S)$ as the space of positive transverse cocycles $\mathcal{H}^+(\lambda)$, in which planes parallel to the boundary are level sets of the hyperbolic length of λ and rays through the origin are Thurston geodesics. Equivalently, if τ is a train track carrying λ then shear coordinates identify $\mathcal{T}(S)$ as a half-space inside $W(\tau)$.

However, shear coordinates are no longer induced by a global projection. Instead, as noted by Thurston, they can be thought of as a map that takes a hyperbolic structure X to (the 1-jet of) its length function with respect to a given lamination. Shear coordinates are then a map not into $W(\tau)$ but into its dual space $W(\tau)^*$ (which can be identified with $W(\tau)$ via the non-degenerate Thurston symplectic form). The image cone is then the positive dual ^[13] of the cone of measures on λ .

This formalism then indicates how shear coordinates generalize to maximal but non-uniquely ergodic laminations. The map is the same, but now the positive dual of $\Delta(\lambda)$ has angles obtained from the intersection of hyperplanes: one for each ergodic measure on λ . Rays in the cone still correspond to geodesics, and affine planes parallel to the bounding planes correspond with the level sets of hyperbolic length of the ergodic measures on λ .

Our shear-shape coordinates come into play when λ is not maximal. In this case, one can go through the above steps for each maximal train track τ , obtained from a snug train track carrying λ by adding finitely many branches. Since λ is carried on a proper subtrack of τ its cone of measures lives in a proper subspace $E \subset W(\tau)$. Taking the positive dual of $\Delta(\lambda)$ and applying the isomorphism $W(\tau) \cong W(\tau)^*$ induced by the Thurston form then realizes Teichmüller space as a cone C in $W(\tau)$. By definition $C \cap E$ is exactly $\mathcal{H}^+(\lambda)$, and one can check this demonstrates C as an affine $\mathcal{H}^+(\lambda)$ bundle.

However, the base of this bundle structure is not canonically determined, in part because $E < W(\tau)$ is generally not symplectic. Moreover, the same hyperbolic structure is parametrized by elements in many different maximal completions, and to achieve $\text{Mod}(S)$ -equivariance one needs to understand how to compare coordinates for different completions. Shear-shape space is designed to solve both of these problems, picking out geometrically meaningful completions and gluing together the corresponding cones all while preserving the bundle structure.

Indeed, the shear-shape coordinates defined in Section ^[13] below associate to each hyperbolic structure a natural finite set of completions (corresponding to standard smoothings of snug train tracks plus geometric arc systems) together with a weight system on each completion. The discussion of this section (Proposition ^[9.5] especially) then implies that the associated shear-shape cocycle is independent of the choice of completion, and that the corresponding train track charts glue together according to the combinatorics of $\mathcal{B}(S \setminus \lambda)$. In this picture, level sets of the hyperbolic length now correspond to bundles over $\mathcal{B}(S \setminus \lambda)$ whose fibers are affine subspaces parallel to the boundary of $\mathcal{H}^+(\lambda)$, while rays in $\mathcal{SH}^+(\lambda)$ correspond to scaling both the coordinate in $\mathcal{B}(S \setminus \lambda)$ as well as the coordinate in $\mathcal{H}^+(\lambda)$.

10. SHEAR-SHAPE COORDINATES FOR TRANSVERSE FOLIATIONS

We now show how the familiar period coordinates for a stratum of quadratic differentials can be reinterpreted as shear-shape coordinates. The main construction of this section is that of the map

$$I_\lambda : \mathcal{F}^{uu}(\lambda) \rightarrow \mathcal{SH}(\lambda)$$

which records the vertical foliation of a quadratic differential and should be thought of as a joint extension of [Mir08, Theorem 6.3] and [MW14, Theorem 1.2].

¹²We remark that this coordinate system is in some sense dual to the paraboloid model of [Thu97, Problem 2.3.13]. Horospherical coordinates place an observer looking out from the center of a family of expanding horospheres, whereas the paraboloid model places an observer at another boundary point looking in.

¹³i.e., those elements of $W(\tau)^*$ which pair positively with every element in $\Delta(\lambda)$ via the intersection form.

The idea is straightforward: given some quadratic differential $q \in \mathcal{F}^{uu}(\lambda)$, the complement $S \setminus Z(q)$ of its zeros deformation retracts onto a neighborhood $N_{\underline{\alpha}(q)}$ of $\lambda \cup \underline{\alpha}(q)$ for some filling arc system $\underline{\alpha}(q)$ (whose topological type reflects the geometry of q). We may therefore identify the period coordinates of q as a relative cohomology class in (the orientation cover of) $N_{\underline{\alpha}(q)}$ with complex coefficients. The imaginary part of this class corresponds to λ , while its real part is the desired shear-shape cocycle $I_\lambda(q)$.

The only obstacle to this plan is in showing that $S \setminus Z(q)$ can actually be identified with a neighborhood of $\lambda \cup \underline{\alpha}(q)$. To overcome this, we recall first in Section 10.1 how to reconstruct the topology of $S \setminus \lambda$ from the horizontal separatrices of q ; this guarantees that all relevant objects have the correct topological types. We then describe in Section 10.2 how to build from $S \setminus Z(q)$ a train track $\tau_{\underline{\alpha}}$ snugly carrying $\lambda \cup \underline{\alpha}(q)$ (Lemma 10.6); this in particular allows us to identify $S \setminus Z(q)$ as a neighborhood of $\lambda \cup \underline{\alpha}(q)$. We may then define $I_\lambda(q)$ using the strategy outlined above and identify it as a weight system on $\tau_{\underline{\alpha}}$ (Lemma 10.10).

Section 10.3 contains a discussion of the global properties of the map I_λ : piecewise linearity, injectivity, and its behavior with respect to the intersection pairing. In this section, we also record Theorem 10.15, which states that I_λ is a homeomorphism onto $\mathcal{SH}^+(\lambda)$. For purposes of convenience, the proof of this theorem is deduced from our later (logically independent) work on shear-shape coordinates for hyperbolic structures (Sections 12–15). See Remark 10.16.

10.1. Separatrices and arc systems. Given a quadratic differential with $|\text{Im}(q)| = \lambda$, our first task towards realizing $|\text{Re}(q)|$ as a shear-shape cocycle is to build a filling arc system $\underline{\alpha}(q)$ on $S \setminus \lambda$ that encodes the horizontal separatrices of q . We begin by recalling how to recover the topology of $S \setminus \lambda$ from the realization of λ as a measured foliation on q .

Recall that a *boundary leaf* ℓ of a component of $S \setminus \lambda$ is a complete geodesic contained in its boundary. Note that infinite boundary leaves of $S \setminus \lambda$ are in 1-to-1 correspondence with leaves of λ which are isolated on one side, while finite boundary leaves (i.e., closed boundary components) are in 2-to-1 correspondence with closed leaves of λ .¹⁴

The corresponding notion for measured foliations is that of *singular leaves*. Let \mathcal{F} be a measured foliation on S and $\tilde{\mathcal{F}}$ denote its full preimage to \tilde{S} under the covering projection; then a bi-infinite geodesic path of horizontal separatrices ℓ is a singular leaf of $\tilde{\mathcal{F}}$ if for every saddle connection s comprising ℓ , the separatrices adjacent to s leave from the same side of ℓ (i.e., always from the left or always from the right); see [Lev83, Figure 2].

There is a fundamental correspondence between boundary leaves of a lamination and singular leaves of a foliation which we record below. Heuristically, collapsing the complementary regions of a lamination yields a foliation; the deflation map of Section 5.3 is a geometric realization of this phenomenon. Again, compare [Lev83, Figure 2] as well as [Min92, Lemma 2.1].

Lemma 10.1. Let λ be a measured lamination on S and let \mathcal{F} be a measure-equivalent measured foliation. Then there is a one-to-one, $\pi_1(S)$ -equivariant correspondence between the boundary leaves of $\tilde{S} \setminus \tilde{\lambda}$ and singular leaves of $\tilde{\mathcal{F}}$. Moreover, singular leaves of $\tilde{\mathcal{F}}$ that share a common separatrix correspond to boundary leaves of the same component of $\tilde{S} \setminus \tilde{\lambda}$.

This lemma in particular allows us to read off the topological type of $S \setminus \lambda$ from the horizontal separatrices of q . Set $\Xi(q)$ to be the union of the horizontal separatrices of q , equipped with the path metric. This 1-complex also comes equipped with a ribbon structure (that is, a cyclic ordering of the edges incident to each vertex) and by thickening each component of $\Xi(q)$ according to this ribbon structure we see that $\Xi(q)$ can be regarded as a spine for the components of $S \setminus \lambda$.

Our construction of $\underline{\alpha}(q)$ then records the dual arc system to the spine $\Xi(q)$ of $S \setminus \lambda$.

Construction 10.2. Let q be a quadratic differential on S with $|\text{Im}(q)| = \lambda$. By the correspondence of Lemma 10.1 each horizontal separatrix of q corresponds to a pair of boundary leaves of the same component of $S \setminus \lambda$. Each infinite separatrix corresponds to a pair of asymptotic boundary leaves, while non-asymptotic boundary leaves are glued along horizontal saddle connections. Dual to each horizontal saddle connection of $\Xi(q)$ is a proper isotopy class of arcs on $S \setminus \lambda$, and we set $\underline{\alpha}(q)$ to be the union of all of these arcs.

Since $\Xi(q)$ is a spine for $S \setminus \lambda$ and $\underline{\alpha}(q)$ consists of arcs dual to its compact edges, we quickly see that

¹⁴This is true because we have insisted that λ support a measure, and so no non-closed leaf may be isolated from both sides.

Lemma 10.3. The arcs of $\underline{\alpha}(q)$ are disjoint and fill $S \setminus \lambda$.

Proof. Each component of $\tilde{S} \setminus \tilde{\lambda}$ has a deformation retract onto the universal cover $\tilde{\Xi}$ of a component of $\Xi(q)$. In particular, as the interiors of the edges of $\tilde{\Xi}$ are disjoint, duality implies that the arcs of $\tilde{\alpha}(q)$ can all be realized disjointly. As this picture is invariant under the covering transformation, this implies that the arcs are disjoint downstairs in $S \setminus \lambda$.

Similar considerations also imply that the arc system is filling: let Σ be a component of $S \setminus \lambda$ with universal cover $\tilde{\Sigma}$ with spine $\tilde{\Xi}$. By construction, the edges of $\tilde{\alpha}(q)$ in $\tilde{\Sigma}$ are dual to the edges of $\tilde{\Xi}$. Since $\Xi(q)$ is a spine for $S \setminus \lambda$, any loop in Σ is homotopic to a union of saddle connections, implying that any nontrivial loop must pass through an edge of $\underline{\alpha}(q)$. Hence $\underline{\alpha}(q)$ fills $S \setminus \lambda$. \square

10.2. Period coordinates as shear-shape cocycles. Now that we understand the relationship between λ and the horizontal data of q , it is easy to build objects $T^* \setminus H^*$ and T^* on q of the same topological type as λ and $\lambda \cup \underline{\alpha}(q)$. However, it is not immediate to actually identify these objects as neighborhoods of λ and $\lambda \cup \underline{\alpha}(q)$. Below, we deduce this from the stronger statement that they admit smoothings onto train tracks snugly carrying λ and $\lambda \cup \underline{\alpha}(q)$; compare [Mir08, Sections 5.2 and 5.3].

Construction 10.4 (Train tracks from triangulations). Let H denote the set of all horizontal saddle connections on q and let T be a triangulation of q containing H . Let T^* be the 1-skeleton of the dual complex to T and let H^* denote the edges of T^* dual to H . Note that T^* is trivalent by definition.

Let Δ denote a triangle of T with dual vertex v_Δ in T^* . Using the $|q|$ -geometry of Δ we may assign tangential data to v_Δ as follows (compare Figures 14 and 15).

- If no edge of Δ is horizontal, then a unique edge e has largest (magnitude of) imaginary part. Assign tangential data to v_Δ so that the dual edge to e is a large half-branch.
- Otherwise, some edge of Δ is horizontal and the other two edges have the same imaginary parts. In this case, we choose tangential data so that the horizontal edge corresponds to a small half-branch and leaves the large half-branch from the right, as seen by the large half-branch.

We denote the resulting train track by $\tau_{\underline{\alpha}}$. The subgraph $T^* \setminus H^*$ can also be converted into a train track τ by deleting the branches of $\tau_{\underline{\alpha}}$ dual to H .

Remark 10.5. We observe that the edges of H^* correspond to the arcs of $\underline{\alpha}(q)$ and $\tau_{\underline{\alpha}}$ is a standard smoothing of $\tau \cup \underline{\alpha}(q)$. Our convention for “standard” ensures that additivity in period coordinates corresponds to additivity in train track coordinates.

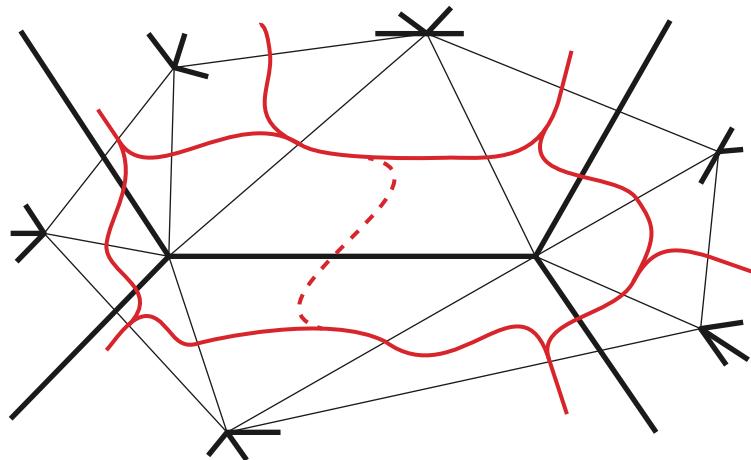


FIGURE 14. An example of the train track $\tau_{\underline{\alpha}}$ around a saddle connection. The thick black lines are stems of horizontal separatrices of q while the light black lines are non-horizontal edges of the triangulation T . The dashed line is a branch of $\tau_{\underline{\alpha}} \setminus \tau$.

By construction, the graph T^* (equivalently, the train track $\tau_{\underline{\alpha}}$) is a deformation retract of $S \setminus Z(q)$. Similarly, $T^* \setminus H^*$ (and τ) are deformation retracts of the complement of the horizontal saddle connections. Together with our discussion above, this implies that τ has the same topological type as λ and $\tau_{\underline{\alpha}}$ has the same topological type as $\lambda \cup \underline{\alpha}(q)$.

In order to actually realize these objects as neighborhoods of λ , we observe that we can build an explicit carrying map from (a foliation measure equivalent to) λ onto τ .

Lemma 10.6. The train track τ carries λ snugly. The weight system on τ that specifies λ is exactly the (magnitude of) the imaginary parts of the periods of the edges of T .

Proof. Let all notation be as above and let \mathcal{F} denote the (singular) horizontal foliation of q .

One can directly build a homotopy of the nonsingular leaves of \mathcal{F} onto τ : in a neighborhood of each edge e of $T \setminus H$ there is a homotopy of the leaves of \mathcal{F} onto the branch of τ dual to e . Now any leaf of \mathcal{F} which passes through a triangle Δ of T does so (locally) only twice and must pass through the side of Δ with the largest imaginary part, which corresponds to a large half-branch of τ . The complement of the separatrix meeting the vertex of Δ opposite to the side with largest imaginary part separates the (locally) non-singular leaves of \mathcal{F} passing through Δ into two packets that can be homotoped onto τ , respecting the smooth structure at the switch dual to Δ ; compare Figure 15.

Now the horizontal foliation \mathcal{F} of q is measure equivalent to λ , and so as τ carries \mathcal{F} it carries λ (snugness follows as τ and λ have the same topological type). The statement about the weight system follows from our description of the carrying map. \square

Now that we have identified τ as a snug train track carrying λ , we may in turn identify a neighborhood of $\lambda \cup \underline{\alpha}(q)$ with (a thickened neighborhood of) $\tau_{\underline{\alpha}}$. With this correspondence established, we may now define $I_{\lambda}(q)$ as the image of the real part of the period coordinates of q under the natural isomorphism on cohomology.

Construction 10.7 (Definition of $I_{\lambda}(q)$). Let $S, \lambda, q, \underline{\alpha}(q)$, and $\tau_{\underline{\alpha}}$ be as above. Set $M_{\underline{\alpha}}$ to be a thickened neighborhood of $\tau_{\underline{\alpha}}$ (in the flat metric defined by q) and let $N_{\underline{\alpha}}$ be a snug neighborhood of $\lambda \cup \underline{\alpha}(q)$ (taken in some auxiliary hyperbolic metric). Perhaps by shrinking $N_{\underline{\alpha}}$, we may assume it embeds into $M_{\underline{\alpha}}$ as a deformation retract (this follows by snugness).

Now $\tau_{\underline{\alpha}}$ is itself a deformation retract of $S \setminus Z(q)$, so the inclusion $M_{\underline{\alpha}} \hookrightarrow S \setminus Z(q)$ is a homotopy equivalence; composing inclusions $N_{\underline{\alpha}} \hookrightarrow M_{\underline{\alpha}} \hookrightarrow S \setminus Z(q)$ and lifting to the orientation covers yields the isomorphism

$$(22) \quad H^1(\widehat{S}, Z(\sqrt{q}); \mathbb{C}) \xrightarrow{j^*} H^1(\widehat{N}_{\underline{\alpha}}, \partial \widehat{N}_{\underline{\alpha}}; \mathbb{C})$$

where the hats denote the corresponding orientation covers. As the composite retraction respects the covering involution ι , this isomorphism also identifies -1 eigenspaces for ι^* . We therefore define

$$I_{\lambda}(q) = \text{Re}(j^* \text{Per}(q))$$

where $\text{Per}(q)$ are the period coordinates for q , and where the real part is taken relative to the natural splitting $\mathbb{C} = \mathbb{R} \oplus i\mathbb{R}$.

Remark 10.8. From the above construction, a basis consisting of branches for the weight space of $\tau_{\underline{\alpha}}$ (equivalently, a basis for $H_1(\widehat{N}_{\underline{\alpha}}, \partial \widehat{N}_{\underline{\alpha}}; \mathbb{Z})$ of dual arcs) picks out a basis for $H_1(\widehat{S}, Z(\sqrt{q}); \mathbb{Z})$. Moreover, each relative cycle is realized geometrically as a saddle connection (as opposed to concatenations, thereof).

To see that $I_{\lambda}(q)$ is indeed a shear-shape cocycle, we need only observe that the values on standard transversals to $\underline{\alpha}(q)$ are all positive. This follows essentially by definition of the orientation cover and construction of $\underline{\alpha}(q)$. To wit: if α is an arc of $\underline{\alpha}(q)$ dual to a saddle connection s , and t is a standard transversal to α , then the canonical lifts of t are mapped to those of s under the isomorphism (22). As the periods of \sqrt{q} increase as you move along the (oriented) horizontal foliation of (\widehat{S}, \sqrt{q}) , this implies that the value of $I_{\lambda}(q)$ on either of the lifts of t is exactly the length of the saddle connection s .

Therefore, we see that the weighted arc system underlying $I_{\lambda}(q)$ is none other than

$$\underline{A}(q) := \sum_{\alpha \in \underline{\alpha}(q)} c_{\alpha} \alpha$$

where c_{α} is the $|q|$ -length of the horizontal saddle connection dual to the arc α .

Remark 10.9. Naturality of all of the isomorphisms involved quickly implies that this construction does not depend on the choice of initial triangulation T . Indeed, suppose that T_1 and T_2 are two triangulations giving rise to train tracks τ_1 and τ_2 and hence shear-shape cocycles σ_1 and σ_2 . Since both τ_i carry $\lambda \cup \underline{\alpha}(q)$ snugly, Lemma 10.6 implies that they have a common refinement τ . Lifting the inclusions

$$N(\tau \cup \underline{\alpha}(q)) \hookrightarrow N(\tau_i \cup \underline{\alpha}(q)) \hookrightarrow S \setminus Z(q)$$

to their orientation covers and drawing the appropriate commutative diagram of cohomology groups, we see that the shear-shape cocycles built from each T_i coincide as weight systems on the common refinement τ .

For use in the sequel, we record below the weight systems on $\tau_{\underline{\alpha}}$ corresponding to λ and $I_{\lambda}(q)$. The proof follows by combining the constructions above with the discussion in Section 9 and is therefore left to the scrupulous reader. See also Figure 15.

For a complex number z , define

$$[z]_+ = \begin{cases} z & \text{if } \arg(z) \in [0, \pi) \\ -z & \text{if } \arg(z) \in [\pi, 2\pi). \end{cases}$$

Observe that $[z]_+ = [-z]_+$ for all $z \in \mathbb{C}$.

Lemma 10.10. Let all notation be as above, and for each edge e of T let b_e denote the branch of $\tau_{\underline{\alpha}}$ dual to it. Then the assignment

$$b_e \mapsto \left[\int_e \sqrt{q} \right]_+$$

defines a complex weight system $w(q)$ on $\tau_{\underline{\alpha}}$ satisfying the switch conditions. Moreover,

$$\text{Im}(w(q)) = \lambda \text{ and } \text{Re}(w(q)) = I_{\lambda}(q).$$

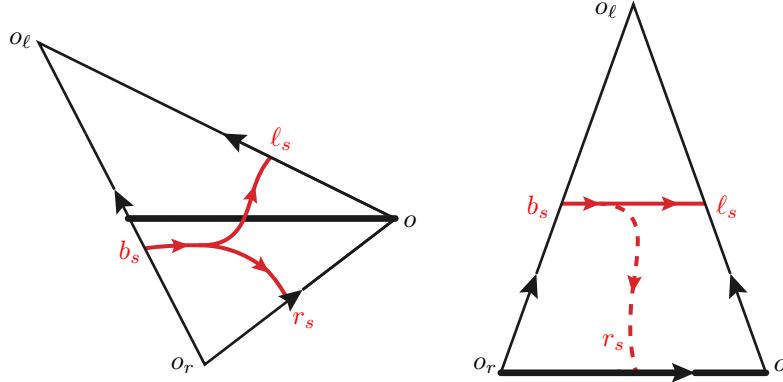


FIGURE 15. Local pictures of the different types of switches of $\tau_{\underline{\alpha}}$. Here we have illustrated the images of each triangle under the holonomy map. The orientation of each edge should be interpreted as indicating the value of $[\cdot]_+$ so that the edge vector is exactly the complex weight assigned to the dual branch of $\tau_{\underline{\alpha}}$. The graphical conventions of this figure mirror those of Figure 14.

10.3. Global properties of the coordinatization. In this section, we show that the map I_{λ} defined above gives a global coordinatization of $\mathcal{MF}(\lambda) \cong \mathcal{F}^{uu}(\lambda)$. First, we record certain global properties of this map; as it is defined by reinterpreting period coordinates as shear-shape cocycles, it preserves many of the structures imposed by period coordinates.

For example, it follows by construction that I_{λ} respects the stratification of each space. That is, if $q \in \mathcal{QT}(k_1, \dots, k_n) \cap \mathcal{F}^{uu}(\lambda)$, then the spine dual to $\underline{\alpha}(q)$ has vertices of valence k_1+2, \dots, k_n+2 . In a similar vein, since both $\mathcal{F}^{uu}(\lambda)$ and $\mathcal{SH}(\lambda)$ have local cohomological coordinates (which induce PIL structures) we can deduce the following:

Lemma 10.11. For any $\lambda \in \mathcal{ML}(S)$, the map I_λ is $\text{Mod}(S)[\lambda]$ -equivariant and PIL.¹⁵

Proof. Equivariance follows from the naturality of our construction: all combinatorial data (arc systems, train tracks, etc.) can be pulled back to a reference surface equipped with λ , so changing the marking by an element of $\text{Mod}(S)[\lambda]$ acts by transforming the combinatorial data on the reference surface.

The piecewise-linear structure on $\mathcal{F}^{uu}(\lambda)$ (respectively, $\mathcal{SH}(\lambda)$) is given by period coordinates (respectively, cohomological coordinates in a neighborhood/train track coordinates) and so the map is by construction piecewise-linear. Integrality comes from the fact that a homotopy equivalence induces an isomorphism on cohomology with \mathbb{Z} -coefficients, hence takes integral points to integral points. \square

The Thurston intersection pairing gives us a powerful tool to understand constraints on the image of I_λ ; in particular, $I_\lambda(q)$ must be a *positive* shear-shape cocycle. Indeed, the tangential structure of the train track τ_α at each switch provides us with an identification of each triangle Δ of T with an oriented simplex. With respect to this orientation, we can compute the area of Δ by taking (one half of) the cross product of two of its sides. Comparing the formula for the cross product with the Thurston intersection pairing (20) then allows us to see that the intersection of λ and $I_\lambda(q)$ is exactly the area of q ; compare [Mir08, Lemma 5.4].

Proposition 10.12. For all $\eta \in \mathcal{MF}(\lambda)$ and all $\mu \in \Delta(\lambda)$,

$$\omega_{\mathcal{SH}}(I_\lambda(\eta), \mu) = i(\eta, \mu).$$

In particular, $I_\lambda(\mathcal{MF}(\lambda)) \subseteq \mathcal{SH}^+(\lambda)$.

The proof of this lemma is made technical by the fact that if μ and $\mu' \in \Delta(\lambda)$ are ergodic but not projectively equivalent then they are mutually singular. To deal with this difficulty, we build a flat structure on the subsurface filled by μ by integrating against $\lambda + t\mu$ and $I_\lambda(\eta)$ for small t . The triangulation T then induces a combinatorially equivalent triangulation of this new flat structure by saddle connections, allowing us to compare the area of this new flat metric (computed via cross products) with the Thurston form on our original train track τ_α . This inverse construction will also be used in the proof of Proposition 10.14.

Proof. We begin by observing that since $\mu \in \Delta(\lambda)$, there is a union of minimal components of the horizontal foliation of $q(\eta, \lambda)$ that supports μ . Call this subfoliation \mathcal{F} and let Y denote the subsurface filled by \mathcal{F} on $q(\eta, \lambda)$. Note that ∂Y must be a union of horizontal saddle connections, hence is contained in any triangulation T used to define τ_α . In particular, $T|_Y$ is a triangulation of Y .

Since η and λ are realized transversely on $q(\eta, \lambda)$ and this specific realization of η is non-atomic (as any closed leaves of η have become vertical cylinders), we can compute the intersection number between η and any measure μ supported on \mathcal{F} as

$$(23) \quad i(\eta, \mu) = \int_S \eta \times \mu = \int_Y \eta \times \mu.$$

We now build a new flat structure on Y whose conical singularities coincide with those of Y ; the salient feature is that $T|_Y$ can be straightened out to a triangulation by saddle connections on the new singular flat structure that reflects the geometry of $\lambda + t\mu$. To construct the new singular flat structure, we build charts from a neighborhood of each triangle $\Delta \subset T|_Y$ to \mathbb{C} and describe the transitions.

Each triangle Δ of T is dual to a switch s with an edge that is dual to a large half-branch b incident to s . Orient $\tau_\alpha \cap \Delta$ so that a train traveling along b toward s is moving in the positive direction. The other edges r and ℓ of Δ are dual to the half-branches of τ_α to the right and left of s , respectively. The vertices o_r, o_ℓ are adjacent to r and ℓ , respectively, and the vertex o is opposite b ; see Figure 15. On the interior of each triangle Δ , we orient the leaves of \mathcal{F} parallel to b . The leaves of η are given the orientation so that the ordered basis of tangent vectors to λ and η at each point agree with the underlying orientation of S . With this orientation, the measures η and λ induce smooth real 1-forms $d\eta$ and $d\lambda$ that look locally like dx and dy , respectively (as opposed to $|dx|$ and $|dy|$, respectively).

Restricted to the interior of Δ , the local orientation of the leaves of η also gives the measure μ the structure of a measurable 1-form that we call $d\mu$. Spreading out the measure on a closed leaf of μ over the horizontal cylinder of λ corresponding to its support as necessary, we get that the map

$$F_t : p \in \Delta \mapsto \int_{\gamma_p} d\eta + id(\lambda + t\mu) \in \mathbb{C}$$

¹⁵We recall that a PL map between PIL manifolds is itself PIL if it sends integral points to integral points.

obtained by integrating along a path γ_p from o_r to p is isometric along leaves of \mathcal{F} and non-decreasing along leaves of η . We compute

$$F_t(o) = I_\lambda(\eta)(r) + i(\lambda + t\mu)(r) \text{ and } F_t(o_\ell) = I_\lambda(\eta)(b) + i(\lambda + t\mu)(b).$$

Transverse invariance and additivity of μ gives

$$(24) \quad F_t(o_\ell) - F_t(o) = I_\lambda(\eta)(\ell) + i(\lambda + t\mu)(\ell).$$

Since the pair $(F_0(o), F_0(o_\ell))$ forms a positively ordered basis for \mathbb{C} (equivalently, since the triangle Δ is positively oriented), the pair $(F_t(o), F_t(o_\ell))$ is also positively oriented for small enough t . Let Δ'_t be the convex hull of $(F_t(o_r), F_t(o), F_t(o_\ell))$.

The area of Δ'_t may now be computed as half the cross product of $F_t(o)$ and $F_t(o_\ell)$. Using equation (24) and linearity of the cross product, we have the formula

$$(25) \quad \text{Area}(\Delta'_t) = \frac{1}{2} \begin{vmatrix} I_\lambda(\eta)(r) & I_\lambda(\eta)(\ell) \\ \lambda + t\mu(r) & \lambda + t\mu(\ell) \end{vmatrix} > 0.$$

Now for each Δ and any small enough t the map F_t may be extended to an open set $U(\Delta)$ in $Y \setminus Z(q)$ that contains Δ (minus its vertices) and so that for every $p \in U(\Delta)$ there is a unique non-singular $|q|$ -geodesic segment γ_p joining o_r to p . We claim that moreover, we may choose $U(\Delta)$ so that $\Delta'_t \subset F(U(\Delta))$; see Figure 16.

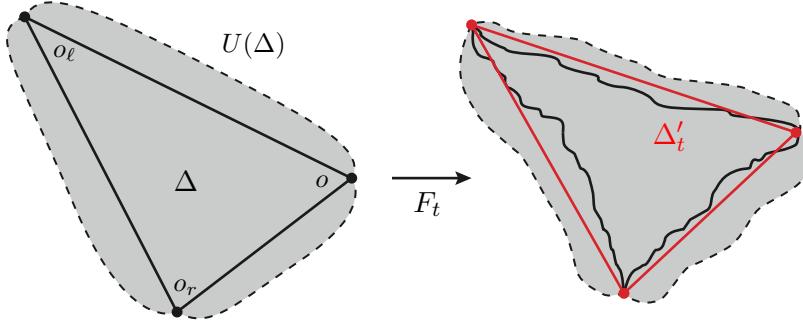


FIGURE 16. Integrating against η and $\lambda + t\mu$ defines a new flat structure on triangles. These charts piece together to give a new half-translation structure on the subsurface filled by μ .

If not, there is some vertex v of $\mathsf{T}_Y \setminus \Delta$ such that $F_t(v) \in \Delta'_t \setminus F_t(\Delta)$. Indeed, by construction, $U(\Delta)$ is a star-shaped neighborhood about the vertex o_r of Δ , so there is a saddle connection joining o_r to v . This saddle connection passes through or shares a vertex of an edge e of Δ . Moreover, we may find v so that the triangle Δ_v formed by e and v is singularity free and contained in $U(\Delta)$. But now, the straightening Δ'_v of $F_t(\Delta_v)$ in \mathbb{C} lies inside Δ'_t with the wrong orientation since $F_t(v)$ lies between $F_t(e)$ and the corresponding edge of Δ'_t . This is a contradiction to the fact that F_t is non-decreasing along leaves of η , alternatively, to the fact that the straightenings Δ'_t are all positively oriented for small enough t . So we may assume that $\Delta'_t \subset F(U(\Delta))$.

If $\Delta_1 \subset \mathsf{T}_Y$ shares an edge with Δ , then the construction of the map F_t on Δ_1 agrees with F_t on $U(\Delta) \cap U(\Delta_1)$ up to multiplication by ± 1 (depending on the configuration of the switches dual to Δ and Δ_1) and translation by the period of the arc connecting the basepoints o_r of each triangle. Thus these triangles glue up to a half-translation structure on $Y \setminus Z$ equipped with a triangulation by saddle connections corresponding to $\mathsf{T}|_Y$.

In our new flat structure on Y , $\lambda + t\mu$ is measure equivalent to the horizontal foliation and (the restriction of) η is equivalent to the vertical foliation. Hence we obtain for any t small enough that

$$\int_Y \eta \times (\lambda + t\mu) = \sum_{\Delta \in \mathsf{T}|_Y} \text{Area}(\Delta'_t) = \sum_{\Delta \in \mathsf{T}|_Y} \frac{1}{2} \begin{vmatrix} I_\lambda(\eta)(r) & I_\lambda(\eta)(\ell) \\ \lambda + t\mu(r) & \lambda + t\mu(\ell) \end{vmatrix} = \omega_{\mathcal{SH}}(I_\lambda(\eta), \lambda + t\mu)$$

where the second equality follows from (25) and the third from (20). Combining this with formula (23) and the linearity of the Thurston intersection form (Lemma 8.3), we get that

$$i(\eta, \mu) = \frac{1}{t} \left(\int_Y \eta \times \lambda + t\mu - \int_Y \eta \times \lambda \right) = \frac{1}{t} (\omega_{\mathcal{SH}}(I_\lambda(\eta), \lambda + t\mu) - \omega_{\mathcal{SH}}(I_\lambda(\eta), \lambda)) = \omega_{\mathcal{SH}}(I_\lambda(\eta), \mu),$$

completing the proof of the lemma. \square

From the proof of Proposition 10.12 we can also extract the following, which allows us to reconstruct a (triangulated) quadratic differential from a sufficiently positive shear-shape cocycle, inverting Construction 10.4.

Lemma 10.13. Let τ be a train track snugly carrying λ and let $\tau_{\underline{\alpha}}$ be a standard smoothing of $\lambda \cup \underline{\alpha}$. Suppose that $\sigma \in \mathcal{SH}(\lambda)$ is represented by a weight system on $\tau_{\underline{\alpha}}$ so at every switch s of $\tau_{\underline{\alpha}}$, the contribution

$$\frac{1}{2} \begin{vmatrix} \sigma(r_s) & \sigma(\ell_s) \\ \lambda(r_s) & \lambda(\ell_s) \end{vmatrix}$$

of s to $\omega_{\mathcal{SH}}(\sigma, \lambda)$ is positive. Then there exists a quadratic differential $q \in \mathcal{F}^{uu}(\lambda)$ so that $I_\lambda(q) = \sigma$ and the dual triangulation to $\tau_{\underline{\alpha}}$ is realized by saddle connections on q .

Proof. The assumption that the contribution at each switch is positive implies that the basis $(F(o), F(o_\ell))$ is positively oriented at each switch, and so we can build a positively-oriented triangle Δ with the prescribed periods. These glue together into the desired quadratic differential. \square

In particular, we can locally invert I_λ by building a quadratic differential out of triangles whose edges have specified periods, so we see that I_λ is injective.

Proposition 10.14. For any $\lambda \in \mathcal{ML}(S)$, the map I_λ is a homeomorphism onto its image.

Proof. To see that I_λ is injective, we observe that Lemma 10.13 provides a (left) inverse map Δ_λ to I_λ . Indeed, suppose that $\sigma = I_\lambda(q)$ for some q and pick a triangulation T as in Construction 10.4, let $\tau_{\underline{\alpha}}$ denote the dual train track. Applying Lemma 10.13 then constructs a quadratic differential q' on which each edge of T is realized as a saddle connection. Since q and q' have the same periods with respect to the same geometric triangulation, they must be equal.

To prove that I_λ is continuous, we first observe that I_λ is by definition continuous on the closure $\mathcal{SH}(\lambda; \underline{\alpha}(q))$ of any cell, as it is induced by a continuous mapping on the level of cohomology. In general, we need only exploit this fact together with a standard reformulation of sequential continuity: a function $f : X \rightarrow Y$ is continuous if and only if every convergent sequence $x_n \rightarrow x$ has a subsequence x_{n_k} so that $f(x_{n_k}) \rightarrow f(x)$.

So let $q_n \rightarrow q \in \mathcal{F}^{uu}(\lambda)$. The polyhedral structure of $\mathcal{SH}(\lambda)$ is locally finite, so for n large enough, $I_\lambda(q_n)$ is contained in a finite union of cells. After passing to a subsequence q_{n_k} , we may assume that q_{n_k} all share the same underlying (maximal) arc system β completing $\underline{\alpha}$. In particular, $I_\lambda(q_{n_k}) \in \mathcal{SH}(\lambda; \beta)$ for all k and so $I_\lambda(q_{n_k}) \rightarrow I_\lambda(q)$ follows from continuity on cells. Therefore I_λ is a continuous injective map between Euclidean spaces of the same dimension (Proposition 8.5 and Corollary 8.2) and so invariance of domain guarantees it is a homeomorphism onto its image. \square

The image of I_λ . In light of Lemma 10.13, to show that I_λ surjects onto $\mathcal{SH}^+(\lambda)$ it would suffice to show that every positive shear-shape cocycle can be realized as a weight system on a train track where every switch contributes positively to the intersection form. However, it is rather complicated to show that every positive shear-shape cocycle admits such a representation (see the discussion in Remark 10.16 just below).

Instead, we deduce this fact using the commutativity of Diagram (2) and the results appearing in Sections 12–15 coordinatizing hyperbolic structures by shear-shape cocycles. We emphasize, however, that Theorem 10.15 is logically independent from the work done in Sections 12–15 that leads to its proof. We include the statement here (as opposed to after Section 15) to provide some closure to our discussion of the parametrization of $\mathcal{MF}(\lambda)$ by shear-shape cocycles.

Theorem 10.15. The map $I_\lambda : \mathcal{F}^{uu}(\lambda) \rightarrow \mathcal{SH}^+(\lambda)$ is a homeomorphism.

Proof. In Section 13 we define the geometric shear-shape cocycle $\sigma_\lambda(X) \in \mathcal{SH}(\lambda)$ associated to a hyperbolic metric $X \in \mathcal{T}(S)$ and show (Theorem 13.13) that $\sigma_\lambda(X) = I_\lambda(\mathcal{O}_\lambda(X))$. In Section 15 we prove Theorem 12.1, which states that the map $\sigma_\lambda : \mathcal{T}(S) \rightarrow \mathcal{SH}^+(\lambda)$ is a homeomorphism. In particular, σ_λ is surjective and hence so is I_λ . Together with Proposition 10.14 this implies the theorem. \square

Remark 10.16. If λ is a maximal lamination, one can deduce surjectivity of I_λ by appealing to the theory of “tangential coordinates” for measured foliations transverse to λ . In general, given τ snugly carrying λ , tangential coordinates can be constructed as a quotient of $\mathbb{R}^{b(\tau)}$ by a vector subspace spanned by vectors that model the change of length of branches of a train track on either side of a switch after a small “fold” or “unzip.” When λ is maximal, there is a linear isomorphism from shear coordinates to tangential coordinates via the symplectic pairing ω_H ; we refer the interested reader to [Thu98, Section 9] or [PH92, §3.4] for details.

The transverse weights defined by the measure of λ on τ together with positive¹⁶ tangential data give τ the structure of a bi-foliated Euclidean band complex. If the tangential data satisfy a collection of triangle-type inequalities, this band complex can be “zipped up” to obtain a bi-foliated flat surface with conical singularities. When defined, the linear transformation mapping tangential coordinates to shear coordinates preserves the intersection number, hence positivity.

A standard positivity argument (see [Thu82, Proposition 9.7.6] or [Thu98, Theorem 9.3]) shows that any tangential data with positive intersection with λ has a positive representative, hence defines a foliation transverse to λ . In particular, the map from $\mathcal{MF}(\lambda)$ to the space of tangential coordinates with positive intersection with λ is surjective. As the space of tangential coordinates with positive intersection is isomorphic to the $\mathcal{H}^+(\lambda)$, this completes the proof of surjectivity in the maximal case.

This being considered, even in the case when λ is maximal “it is harder to see the [positivity] inequalities satisfied by the shear coordinates [than the tangential coordinates]” [Thu98, p. 45] and it is not clear how to run the “standard positivity argument” without passing through tangential coordinates. We have therefore chosen to prove Theorem 10.15 in a way that avoids developing a theory of tangential coordinates dual to shear-shape cocycles. Instead, we take advantage of the relationship between the Thurston intersection form on $\mathcal{SH}(\lambda)$ and the length of λ on a given hyperbolic surface, as exploited in the proof of Theorem 12.1 (see in particular Claim 15.8).

11. FLAT DEFORMATIONS IN SHEAR-SHAPE COORDINATES

The identification of Section 10 between periods of saddle connections and the values of the shear-shape cocycle $I_\lambda(q)$ immediately allows us to transport certain flows on $\mathcal{F}^{uu}(\lambda)$ to shear-shape space. Moreover, Theorem 10.15 affords a new perspective on the “tremor deformations” of [CSW20] (see Definition 11.3).

The horizontal stretch. We begin by observing that the space $\mathcal{SH}^+(\lambda)$ carries a natural $\mathbb{R}_{>0}$ action given by scaling both the underlying arc system \mathcal{A} and the values assigned to test arcs (equivalently, the corresponding cohomology class or the weights on a train track realization). Using our correspondence between period coordinates and shear-shape cocycles (Lemma 10.10), we see that this dilation expands the real part of each period, so the corresponding flat deformation is just a horizontal stretch.¹⁷

Lemma 11.1. Let $q \in \mathcal{F}^{uu}(\lambda)$; then

$$(26) \quad I_\lambda \left(\begin{pmatrix} e^t & 0 \\ 0 & 1 \end{pmatrix} q \right) = e^t I_\lambda(q)$$

for all $t \in \mathbb{R}$.

In particular, we see that our coordinatization linearizes the expansion of the strong unstable foliation under the Teichmüller geodesic flow.

Horocycle flow and tremors. We now consider the horocycle flow on $\mathcal{F}^{uu}(\lambda)$, which is just the restriction of the standard horocycle flow h_s to the strong unstable leaf. An easy computation shows that for every

¹⁶Here, positive means that there is a representative of the tangential data that is positive on each branch of τ .

¹⁷This is just the Teichmüller geodesic flow normalized so that the horizontal foliation remains constant. Applying the standard geodesic flow takes $(I_\lambda(q), \lambda)$ to $(e^{t/2} I_\lambda(q), e^{-t/2} \lambda)$.

saddle connection e of q , one has

$$(27) \quad \left[\int_e \sqrt{h_s q} \right]_+ = \left(\operatorname{Re} \left[\int_e \sqrt{q} \right]_+ + s \operatorname{Im} \left[\int_e \sqrt{q} \right]_+ \right) + i \operatorname{Im} \left[\int_e \sqrt{q} \right]_+$$

(here we have invoked the $[\cdot]_+$ function to avoid fussing over square roots and orientations).

With the help of Lemma 10.10 we may translate this into the language of transverse and shear-shape cocycles to observe

Lemma 11.2. The map I_λ takes horocycle flow to translation by λ in a time preserving way. In symbols,

$$I_\lambda(h_s q) = I_\lambda(q) + s\lambda.$$

More generally, we can perform a similar deformation for *any* measure μ supported on λ , resulting in the *tremor flow* along μ . First defined by Chaika, Smillie, and Weiss in the context of Abelian differentials, the *tremor* $\operatorname{trem}_\mu(q)$ of a quadratic differential $q = q(\eta, \lambda)$ by a measure $\mu \in \Delta(\lambda)$ is the unique quadratic differential specified by shearing η by μ and leaving λ fixed. Why this makes sense (note that η and μ may not fill S) and why it can be continued for all time present significant technical challenges in [CSW20] (see §§4 and 13 therein). However, when considered in our coordinates (and restricted to a leaf of the unstable foliation) tremors become quite simple.

For a given lamination λ , let $|\Delta(\lambda)|_\pm$ denote the vector space of all signed transverse measures on λ ; this is naturally a vector subspace of $\mathcal{H}(\lambda)$ of dimension at most $3g - 3$ with basis consisting of the length 1 (with respect to some auxiliary hyperbolic metric) ergodic measures on λ .

Definition 11.3. Let $q \in \mathcal{F}^{uu}(\lambda)$ and let $\mu \in |\Delta(\lambda)|_\pm$. Then the tremor $\operatorname{trem}_\mu(q)$ of q along μ is the unique quadratic differential specified by

$$(28) \quad I_\lambda(\operatorname{trem}_\mu(q)) = I_\lambda(q) + \mu.$$

Note that the fact that $I_\lambda(q) + \mu \in \mathcal{SH}^+(\lambda)$ follows by affinity of the Thurston form (Lemma 8.3).

Remark 11.4. Technically, the deformation considered above is a “non-atomic tremor” in the language of [CSW20]. One can also consider “atomic tremors,” which transform q by twisting along certain admissible loops of horizontal saddle connections.

In shear-shape coordinates, these admissible loops correspond to certain simple closed curves in the complementary subsurfaces. Atomic tremors are then realized by appropriately shearing the underlying arc system $\underline{A}(q)$ along the curves and transporting the transverse cocycle using the affine connection coming from train-track coordinates. Of course, one can also define tremors along more complicated laminations contained in $S \setminus \lambda$ as well.

For the convenience of the reader familiar with the terminology of [CSW20], we have included a dictionary which translates between our notation and theirs (at least when the horizontal lamination is filling — when it is not, one must replace $\Delta(\lambda)$ with a subset of the zero set of λ and take more care). See Figure 17.

Shear-shape cocycles	Foliation cocycles
$\Delta(\lambda)$	C_q^+
$ \Delta(\lambda) _\pm$	\mathcal{T}_q
$\omega_{\mathcal{SH}}(\eta, \mu) = i(\mu_+, \eta) - i(\mu_-, \eta)$	signed mass $L_q(\mu)$
$i(\mu_+, \eta) + i(\mu_-, \eta)$	total variation $ L _q(\mu)$

FIGURE 17. Translating between our language of shear-shape cocycles and the “foliation cocycles” of [CSW20]. Throughout, we assume that $q = q(\eta, \lambda)$ where λ is filling (equivalently, q has no loops of horizontal saddle connections). We have written a signed transverse measure μ as $\mu = \mu_+ - \mu_- \in |\Delta(\lambda)|_\pm$, where $\mu_\pm \in \Delta(\lambda)$.

We can now immediately deduce certain properties of the tremor map from the structure of $\mathcal{SH}^+(\lambda)$ and the intersection pairing. While we will not use these results in the sequel, we have chosen to include them

in order to demonstrate the utility of our new perspective on these deformations. For example, using our coordinates one can easily deduce that (non-atomic) tremors leave horizontal data invariant and hence can be continued indefinitely while remaining in the same stratum.

Lemma 11.5. For any $q \in \mathcal{F}^{uu}(\lambda)$ and $\mu \in |\Delta(\lambda)|_{\pm}$, the tremor path $\text{trem}_{t\mu}(q)$ is defined for all time and is completely contained in $\mathcal{SH}(\lambda; \underline{A}(q))$. In particular, $\{\text{trem}_{t\mu}(q)\}$ always remains in the same stratum.

Remark 11.6. The above Lemma is one specific instance of a much more general phenomenon. The global description of $\mathcal{F}^{uu}(\lambda)$ afforded by shear-shape coordinates allows one to formulate a general criterion for extending affine period geodesics, a topic which the authors hope to address in future work.

Using our interpretation of tremors as translation, it is similarly easy to describe how tremors interact with other flat deformations. Compare with Propositions 6.1 and 6.5 of [CSW20]. We leave proofs to the reader, as they follow immediately from (28) and (26).

Lemma 11.7. Let $q \in \mathcal{F}^{uu}(\lambda)$. Then for any $\mu \in |\Delta(\lambda)|_{\pm}$ and for $g_t = \begin{pmatrix} e^{t/2} & 0 \\ 0 & e^{-t/2} \end{pmatrix}$, we have that

$$g_t \text{trem}_{\mu}(q) = \text{trem}_{e^{t/2}\mu}(g_t(q)).$$

Additionally, for any $\mu_1, \mu_2 \in |\Delta(\lambda)|_{\pm}$, we have that

$$\text{trem}_{\mu_1}(q) \text{trem}_{\mu_2}(q) = \text{trem}_{\mu_1 + \mu_2}(q) = \text{trem}_{\mu_2}(q) \text{trem}_{\mu_1}(q).$$

In particular, tremors commute with the horocycle flow.

12. SHEAR-SHAPE COORDINATES FOR HYPERBOLIC METRICS

We now parametrize hyperbolic structures on S by shear-shape cocycles for a measured geodesic lamination λ . With respect to the Lebesgue measure on $\mathcal{ML}(S)$, the generic lamination cuts a hyperbolic surface into ideal triangles. As all ideal triangles are isometric, Bonahon and Thurston's shearing coordinates need only take into account the "shear" between pairs of complementary triangles to describe a hyperbolic structure. As our objective is to generalize these coordinates to laminations with arbitrary topology, we must therefore combine the data of the geometry of hyperbolic metrics in complementary subsurfaces with the shearing data between them. Shear-shape space $\mathcal{SH}(\lambda)$ is well suited to this task.

In the following Sections 13–15, we explain how to associate a "geometric shear-shape cocycle" to a hyperbolic metric and prove that the space of positive shear-shape cocycles coordinatizes Teichmüller space:

Theorem 12.1. The map $\sigma_{\lambda} : \mathcal{T}(S) \rightarrow \mathcal{SH}^{+}(\lambda)$ that associates to a hyperbolic metric its geometric shear-shape cocycle is a stratified real-analytic homeomorphism.

As detailed in the Introduction, combining this theorem with Theorems 2.1 and 10.15 implies that the orthogeodesic foliation map \mathcal{O}_{λ} is a homeomorphism, and consideration of the earthquake/horocycle flows in $\mathcal{SH}^{+}(\lambda)$ coordinates then proves the conjugacy on slices (Theorem D).

We remark that the stratified regularity of σ_{λ} and \mathcal{O}_{λ} is the best one can expect, since the adjacency of strata of differentials is not analytic (as there are multiple inequivalent ways to "break up a zero"). Compare with [Dum15, Theorem D], in which it is shown that for a fixed Riemann surface Z , the identification $Q(Z) \cong \mathcal{ML}$ guaranteed by the Hubbard-Masur theorem [HM79] is stratified real-analytic.

Fixed complementary subsurfaces. By definition (see Section 13.2), the weighted arc system $\underline{A}(X)$ underlying $\sigma_{\lambda}(X)$ exactly identifies the geometry of $X \setminus \lambda$ via Theorem 6.4. Setting

$$\mathcal{T}(S; \underline{A}) := \{X \in \mathcal{T}(S) : \underline{A}(X) = \underline{A}\},$$

Theorem 12.1 therefore implies that $\mathcal{T}(S; \underline{A})$ is nonempty if and only if $\underline{A} \in \mathcal{B}(S \setminus \lambda)$.

Remark 12.2. The authors do not know a proof of this fact that does not factor through Theorem 12.1 except in some special cases (for example, when the complement of λ is polygonal, or when λ is a union of simple closed curves). ¹⁸

¹⁸One can of course complete λ to a maximal lamination and then specify the shear coordinates on each of the added leaves, but then one must be very careful to ensure that these shears satisfy the relations coming from the metric residue condition. The argument then requires an involved computation with train tracks carrying the completed lamination.

In fact, since $\mathcal{SH}^+(\lambda)$ is an affine cone bundle over $\mathcal{B}(S \setminus \lambda)$ (Proposition 8.5), we see that

Corollary 12.3. For each $\underline{A} \in \mathcal{B}(S \setminus \lambda)$, the set $\mathcal{T}(S; \underline{A})$ is a real-analytic submanifold of $\mathcal{T}(S)$ and the restriction of σ_λ to

$$\mathcal{T}(S; \underline{A}) \rightarrow \mathcal{SH}^+(\lambda; \underline{A}) \cong \mathcal{H}^+(\lambda)$$

is a real-analytic homeomorphism.

In this setting, the correspondence between $\mathcal{T}(S; \underline{A})$ and $\mathcal{H}^+(\lambda)$ is a natural generalization of shear coordinates, since the complementary subsurfaces to λ are always isometric. In fact, the shape-shifting deformations built to deform X by some $\mathfrak{s} \in \mathcal{H}(\lambda)$ (see the proof sketch of Theorem 12.1 just below) restrict to cataclysms/shear maps in the sense of [Bon96, Section 5]. In particular, if \mathfrak{s} represents a measure supported on λ , then the shape-shifting deformation determined by \mathfrak{s} is part of an earthquake in \mathfrak{s} (Corollary 15.2); if \mathfrak{s} is a multiple of $\sigma_\lambda(X)$, the shape-shifting transformation can sometimes be identified with part of a (generalized) stretch ray (Propositions 15.12 and 15.18).

In addition to being non-empty, $\mathcal{T}(S; \underline{A})$ is structurally rich; the authors hope to explore this space further in future work. Of particular interest is the (degenerate) Weil–Petersson pairing on this locus and its relation with the Thurston symplectic form and Masur–Veech measures.

A sketch of the proof. Since the proof of Theorem 12.1 spans several sections (two of which consist of involved constructions of the relevant objects), we devote the remainder of this section to a broad-strokes outline of the arguments involved. Our exposition throughout these sections is mostly self-contained, but we sometimes refer to [Bon96] for proofs and to [Thu98] for inspiration.

We begin in Section 13 by defining the map σ_λ . Under the correspondence established in Theorem 6.4, we associate to X the weighted arc system $\underline{A}(X)$ recording the hyperbolic structure on $X \setminus \lambda$. We cut X along the (ortho)geodesic realization of $\lambda \cup \underline{\alpha}$ into a union of (degenerate) right-angled polygons, and measure the shear between certain pairs of polygons. We then argue using train tracks that it suffices to record the shearing data comprising $\sigma_\lambda(X)$ on short enough arcs k transverse to λ and disjoint from $\underline{\alpha}(X)$. The value of $\sigma_\lambda(X)$ on short k may then be defined by isotoping k to a path connecting vertices of the spine \mathbf{Sp} and built of segments alternating between leaves of λ and of $\mathcal{O}_\lambda(X)$, then measuring the total (signed) length along λ . These measurements are equivalent to Bonahon and Thurston’s method of measuring shears (via the horocyclic foliation) when k is short enough, but cannot be globally derived from theirs due to obstructions coming from complementary subsurfaces.

The proof that σ_λ is a homeomorphism then follows the same general steps as appear in [Bon96]. After proving that σ_λ is injective and lands inside of $\mathcal{SH}^+(\lambda)$ (Proposition 13.12 and Corollary 13.14), we then show that it is open (Theorem 15.1) and proper. Since $\mathcal{SH}^+(\lambda)$ is a cell (Proposition 8.5), invariance of domain then implies that σ_λ must be a homeomorphism.

Our proof of injectivity mirrors that of [Bon96, Theorem 12] with an additional invocation of Theorem 6.4. For properness we mostly appeal to [Bon96, Theorem 20] but need to discuss complications that arise from the piecewise-linear structure of shear-shape space. Similarly, our broad-strokes strategy to prove openness parallels that of [Bon96, §5], in that we build a “shape-shifting cocycle” $\varphi_\mathfrak{s}$ for all small enough deformations \mathfrak{s} of $\sigma_\lambda(X)$ (see Section 14). Deforming X by post-composing its charts to \mathbb{H}^2 with $\varphi_\mathfrak{s}$ then yields a surface $X_\mathfrak{s}$ with $\sigma_\lambda(X_\mathfrak{s}) = \sigma_\lambda(X) + \mathfrak{s}$.

It is in the construction of $\varphi_\mathfrak{s}$, performed in Section 14, where our discussion truly diverges from [Bon96] and [Thu98]. When λ is maximal, one can specify $\varphi_\mathfrak{s}$ by shearing X along the leaves of λ (i.e., performing a cataclysm). Even in the maximal case this procedure is delicate, hinging on the convergence of infinite products of small Möbius transformations (compare Section 14.2). In the non-maximal case, we must also simultaneously account for the changing shapes of complementary subsurfaces (which also introduces extra complications into the shearing deformations since the shapes of spikes are changing). See the introduction to Section 14 for a more granular description of the construction of $\varphi_\mathfrak{s}$.

13. MEASURING HYPERBOLIC SHEARS AND SHAPES

In this section, we take our first steps towards proving Theorem 12.1 by describing how to associate to any hyperbolic surface X a *geometric shear-shape cocycle* $\sigma_\lambda(X)$ in a natural way; this yields the map

$$\sigma_\lambda : \mathcal{T}(S) \rightarrow \mathcal{SH}(\lambda).$$

After fixing some notational conventions that we will use throughout the sequel, we define $\sigma_\lambda(X)$ by first specifying its underling arc system $\underline{A}(X)$ in a variety of equivalent ways. After doing so, we define the shear between “nearby” hexagons analogously to Bonahon and Thurston; placing all of this data onto a standard smoothing $\tau_{\underline{\alpha}}$ of a geometric train track is therefore enough to specify $\sigma_\lambda(X)$ (Lemma 13.6).

We then show that the data of shears between any two nearby hexagons can be recovered from the weight system on $\tau_{\underline{\alpha}}$, even if those hexagons are not “visible” to $\tau_{\underline{\alpha}}$ (Lemma 13.9). This in particular implies that our choice of $\tau_{\underline{\alpha}}$ does not actually matter, and hence $\sigma_\lambda(X)$ is well-defined.

We then conclude the section by proving some initial properties of σ_λ . Proposition 13.12 shows that the map is injective following an argument of Bonahon, while in Theorem 13.13 we show that our map captures the geometry of the orthogeodesic foliation.

13.1. Preliminaries and notation. In this section, we discuss the geometry of a geodesic lamination on a hyperbolic surface and fix notation in preparation of our definition of the geometric shear-shape cocycle of a hyperbolic structure.

Throughout, we use the symbol λ to refer to both the measured lamination λ and its support, realized geodesically with respect to any number of hyperbolic metrics. We refer to Remark 7.10 for a discussion of how to relax the assumption that λ is measured. We reserve S to denote a topological surface and Σ the topological type of a component of $S \setminus \lambda$, while X and Y will denote their hyperbolic incarnations. We also adopt the following family of notational conventions: the expression $g \subset \lambda$ means that g is a leaf of λ , and $Y \subset X \setminus \lambda$ means that Y is a component of (the metric completion of) $X \setminus \lambda$, etc. The notation of [Bon96] is used as inspiration, since we will make direct appeals to the results therein. However, our situation requires more care, since we have more objects to keep track of. A key difference is that we will focus not on the relative shear between complementary subsurfaces of $X \setminus \lambda$, but on the relative positioning of pairs of boundary leaves of λ , equipped with a natural collection of basepoints determined by the orthogeodesic foliation.

Hexagons. Given $X \in \mathcal{T}(S)$ and $\lambda \in \mathcal{ML}(S)$, realize λ geodesically on X . Construct the orthogeodesic foliation $\mathcal{O}_\lambda(X)$ on X with piecewise geodesic spine \mathbf{Sp} and dual arc system $\underline{\alpha} = \underline{\alpha}(X)$, realized orthogeodesically with respect to X and λ . The union $\lambda_{\underline{\alpha}} = \lambda \cup \underline{\alpha}$ is a geometric object on X that fills; that is, the metric completion of $X \setminus \lambda_{\underline{\alpha}}$ is a union of geometric pieces that are topological disks, possibly with some points on the boundary removed corresponding to spikes. We lift the situation to universal covers $\widetilde{\lambda}_{\underline{\alpha}} \subset \widetilde{X}$, where we have also the full preimages $\widetilde{\mathbf{Sp}}, \widetilde{\lambda}, \widetilde{\underline{\alpha}}$, etc., of various objects.

Let \mathcal{H} be the vertex set of \mathbf{Sp} ; we will sometimes refer to $v \in \mathcal{H}$ as a *hexagon*. Indeed, to v there is associated a component H_v of $\widetilde{X} \setminus \widetilde{\lambda}_{\underline{\alpha}}$ which is generically a degenerate right-angled hexagon, though H_v may also be a regular ideal or right-angled polygon, for example. We reiterate that, by abuse of terminology, *any complementary component H_v of $\widetilde{X} \setminus \widetilde{\lambda}_{\underline{\alpha}}$ is called a hexagon*, no matter its shape.

If $\{H_v : v \in \mathcal{H}\}$ contains components that are not degenerate right-angled hexagons in the usual sense, then $\underline{\alpha}$ corresponds to a simplex of $\mathcal{A}_{\text{fill}}(S \setminus \lambda)$ of non-maximal dimension (or the empty set, if λ is filling and $\underline{\alpha}$ is empty). One may always include $\underline{\alpha}$ in a maximal arc system $\underline{\beta}$, which necessarily defines a simplex of full dimension. The complementary components of $\widetilde{X} \setminus \widetilde{\lambda}_{\underline{\beta}}$ are now degenerate right-angled hexagons in the usual sense, and gluing them in pairs along $\underline{\beta} \setminus \underline{\alpha}$ gives the more general “hexagons” of $\widetilde{X} \setminus \widetilde{\lambda}_{\underline{\alpha}}$. We will often tacitly choose and work with a maximal arc system containing the original when convenient.

Pointed geodesics. We now define a natural family of basepoints associated to boundary leaves of $\widetilde{\lambda}$. For $v \in \mathcal{H}$ and its associated hexagon H_v , define the λ -boundary $\partial_\lambda H_v$ of H_v to be the set of leaves of $\widetilde{\lambda}$ that meet ∂H_v .

For $v \in \mathcal{H}$ and g a leaf of $\partial_\lambda H_v$, define p_v to be the orthogonal projection of v to g . Observe that v and p_v lie along the same (singular) leaf of $\mathcal{O}_\lambda(X)$. The orientation of S gives H_v an orientation and hence orients ∂H_v ; this yields an orientation-preserving, isometric identification of (g, p_v) with $(\mathbb{R}, 0)$. We refer to points on a based geodesic by their signed distance to the basepoint, so that 0 refers to p_v while $\pm x$ refer to the points at signed distance $\pm x$ from p_v .

For a pair $v \neq w \in \mathcal{H}$ not in the same component of $\widetilde{\mathbf{Sp}}$, there is a unique geodesic $g_v^w \in \partial_\lambda H_v$ that separates v from w . Symmetrically, there is such a pointed geodesic $g_w^v \in \partial_\lambda H_w$ separating w from v . Note

that $g_v^w = g_w^v$ occurs if and only if this leaf is isolated, and by the assumption that λ is measured, projects to a simple closed curve component of λ . Even in this case, the points p_v and p_w are in general different.

13.2. The shear-shape cocycle of a hyperbolic structure. Our first task towards defining the geometric shear-shape cocycle $\sigma_\lambda(X)$ of a hyperbolic structure X is to construct a weighted, filling arc system $\underline{A}(X) \in \mathcal{B}(S \setminus \lambda)$ which records the shapes of the complementary subsurfaces.

With the technology we have developed up to this point, we now have many ways of constructing $\underline{A}(X)$, all of which are easily seen to be equivalent:

- To each $\alpha \in \underline{\alpha}(X)$, we associate the weight $c_\alpha := i(\mathcal{O}_\lambda(X), e_\alpha)$, where e_α is the edge of Sp dual to α . Equivalently, c_α is the length of the projection of e_α to either of the two leaves of λ to which it is closest. Then set $\underline{A}(X) = \sum c_\alpha \alpha$.
- Each component $Y \subset X \setminus \lambda$ is naturally endowed with a hyperbolic structure; by Theorem 6.4 this metric corresponds to a weighted, filling arc system in $|\mathcal{A}_{\text{fill}}(Y, \partial Y)|_{\mathbb{R}}$, and we let $\underline{A}(X)$ denote the union of these arc systems over all components of $X \setminus \lambda$.
- Let q be the quadratic differential with $|\text{Re}(q)| = \mathcal{O}_\lambda(X)$ and $|\text{Im}(q)| = \lambda$; then set $\underline{A}(X) = \underline{A}(q)$.

The final definition together with the results of Section 10 implies that $\underline{A}(X) \in \mathcal{B}(S \setminus \lambda)$ for every hyperbolic structure X on S . In the interest of providing the reader with geometric intuition for this condition, we have included an alternative, purely hyperbolic-geometric proof of this fact below.

Lemma 13.1. With notation as above, $\underline{A}(X) \in \mathcal{B}(S \setminus \lambda)$.

Proof. By Theorem 6.4, it suffices to show that for each minimal, orientable component μ of λ , the sum of the metric residues of the crown ends of $X \setminus \lambda$ incident to μ is 0. If μ is a simple closed curve, then the metric residue is just equal to the (signed) lengths of the boundary components resulting from cutting along μ , which clearly must match.

So assume that μ is not a closed curve and pick an orientation. Construct a geometric train track τ snugly carrying μ as in Construction 5.6; then τ inherits an orientation from the inclusion of μ and so has well defined left- and right-hand sides. As in Section 5.2, every branch b of τ has a well-defined length along λ which we denote by $\ell_\tau(b) > 0$. At each switch s of τ , let h_s be the leaf of the horocyclic foliation of $\mathcal{N}_\epsilon(\mu)$ projecting to s . By assumption of snugness, the spikes of $S \setminus \tau$ correspond with the spikes of $S \setminus \mu$, so the union of the h_s truncates each spike of each crown end incident to μ by h_s .

Each crown incident to μ inherits an orientation from the chosen orientation on μ , and we now compute the total metric residue with respect to these orientations and the truncations induced by the h_s 's. Recall that the metric residue of an oriented crown \mathcal{C} is the alternating sum of the lengths of the geodesic boundary segments running between the truncation horospheres (Definition 4.3). Each such geodesic segment defines a co-oriented trainpath $(b_1 \cdot \dots \cdot b_n, \pm)$ in τ (i.e., a trainpath and a distinguished side, left or right, corresponding to $+$ and $-$, respectively) which runs along the entirety of a smooth component of the boundary of $X \setminus \tau$. Using this identification, we may compute that the corresponding contribution to the total metric residue is given by $\pm \sum_i \ell_\tau(b_i)$.

Finally, we observe that every branch of τ is a subpath of exactly two smooth boundary edges of $X \setminus \tau$ (corresponding to its left and right sides). Therefore, the sum of the metric residues of all of the crown ends incident to μ is the sum of the contributions of the corresponding co-oriented trainpaths, which is necessarily 0 since each branch is counted twice, once with positive and once with negative sign. Thus $\underline{A}(X) \in \mathcal{B}(S \setminus \lambda)$. \square

Shears between nearby hexagons. Our second step towards defining $\sigma_\lambda(X)$ is to determine how to record shearing data between two hexagons that lie in different components of $\tilde{X} \setminus \tilde{\lambda}$ yet are close enough together. Except for sign conventions (see Remark 13.3), our discussion is essentially identical to Bonahon's definition of shearing between the plaques of a maximal lamination [Bon96, §2]. Our restriction to pairs of nearby hexagons reflects the fact that if two hexagons are far apart, a path connecting them may meet a subsurface of $\tilde{X} \setminus \tilde{\lambda}$ in a variety of ways.

Given $v, w \in \mathcal{H}$, consider the associated pointed geodesics $(g_v^w, p_v) \in \partial_\lambda H_v$ closest to H_w and $(g_w^v, p_w) \in \partial_\lambda H_w$ closest to H_v . We say that the geodesic segment $k_{v,w} \subset \tilde{X}$ joining p_v to p_w is a *simple piece* if $k_{v,w}$ projects to a simple geodesic segment in X and $k_{v,w}$ bounds a spike in every hexagon that it crosses. That is, if $k_{v,w}$ crosses H_u for some $u \in \mathcal{H}$, then $k_{v,w} \cap H_u$ bounds a triangle in H_u , two sides of which lie on

asymptotic leaves g_v^v and g_w^w defining a spike of $\tilde{\lambda}$. If $k_{v,w}$ is a simple piece, then we say that (v, w) is a *simple pair*.

We observe that if $v, w \in \mathcal{H}$ are close enough together and lie in different components of $\widetilde{\mathbf{Sp}}$, then (v, w) is a simple pair. The exact value of “close enough” is unimportant, but we note that it suffices for $d(H_u, H_v)$ to be smaller than the length of the shortest arc of $\underline{\alpha}(X)$.

Now following Bonahon [Bon96, Section 2], let $\Lambda_{v,w}$ be the leaves of $\tilde{\lambda}$ that separate g_v^w from g_w^v , equipped with the linear order $<$ induced by traversing $k_{v,w}$ from p_v to p_w . Since (v, w) is a simple pair, the subset of those leaves that are also the boundary of a complementary component of $\tilde{X} \setminus \tilde{\lambda}$ come in pairs that are asymptotic in one direction. The partial horocyclic foliations on the wedges bounded by pairs of asymptotic boundary leaves extend across the leaves of $\Lambda_{v,w}$, foliating the region bounded by g_v^w and g_w^v . In particular, the leaf of the horocyclic foliation containing p_v meets g_w^v (and the leaf containing p_w meets g_v^w).

Since the orthogeodesic foliation is equivalent to the horocyclic foliation in spikes, we see that if (v, w) is a simple pair then the leaf of $\mathcal{O}_{\Lambda_{v,w}}(\tilde{X})$ containing p_v meets g_w^v (and the leaf containing p_w meets g_v^w). In fact, simplicity implies that $\mathcal{O}_{\Lambda_{v,w}}(\tilde{X})$ foliates the “quadrilateral” bounded by g_w^v , g_v^v , and the two leaves of $\mathcal{O}_{\Lambda_{v,w}}(\tilde{X})$ containing p_v and p_w .

Definition 13.2. Suppose that (v, w) is a simple pair of hexagons. Using the orientation conventions of Subsection 13.1, identify the corresponding pointed geodesics (g_v^w, p_v) and (g_w^v, p_w) with $(\mathbb{R}, 0)$. Now since the hexagons are close enough, the singular leaf of $\mathcal{O}_{\Lambda_{v,w}}(\tilde{X})$ containing p_v meets g_w^v in some point $r \in \mathbb{R}$, and we set $\sigma_\lambda(X)(v, w) = -r$. See Figure 18.

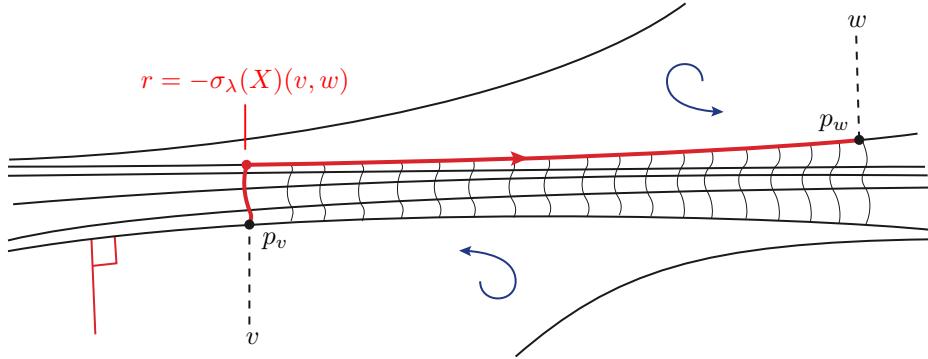


FIGURE 18. Computing the shears between two nearby hexagons v and w . In this example, $r < 0$, so $\sigma_\lambda(X)(v, w) > 0$.

It is not hard to see that $\sigma_\lambda(X)(v, w)$ remains the same if we flip the roles of v and w . Indeed, following along the leaves of the orthogeodesic foliation defines an orientation reversing isometry from a subsegment of g_v^w to a subsegment of g_w^v that takes $t \mapsto r - t$. In particular, p_v maps to a point on g_w^v that is positioned r signed units away from p_w , and so we see that $\sigma_\lambda(X)(v, w) = \sigma_\lambda(X)(w, v)$.

Remark 13.3. Our choice to set $\sigma_\lambda(X)(v, w) = -r$ instead of $+r$ records “how far along g_w^v you must travel from r to get to p_w .” Though this convention is the opposite of what appears in [Bon96], it allows us to combine the data of $\sigma_\lambda(X)(v, w)$ and $\underline{A}(X)$ into a system of train track weights on a standard smoothing (Construction 13.5 just below). Our convention also parallels our choice of $[\cdot]_+$ function when measuring periods of a quadratic differential (Lemma 10.10), which makes the relationship between the hyperbolic geometry of (X, λ) and the flat geometry of $q(\mathcal{O}_\lambda(X), \lambda)$ more transparent.

Below, we give an elementary estimate that will be used in the proof of Proposition 13.12, compare with [Bon96, Lemma 8].

Lemma 13.4. Suppose that (v, w) is a simple pair of hexagons. Let (g_v^w, p_v) and (g_w^v, p_w) be the associated pointed geodesics. Then the geodesic segment $k_{v,w}$ joining p_v to p_w satisfies

$$|\sigma_\lambda(X)(k_{v,w})| \leq \ell(k_{v,w}).$$

Proof. As (v, w) is simple, the partial orthogeodesic foliation $\mathcal{O}_{\Lambda_{v,w}}(\tilde{X})$ foliates the region U bounded by g_w^v , g_v^w , and the two leaves of $\mathcal{O}_{\Lambda_{v,w}}(\tilde{X})$ containing p_v and p_w . This foliation gives rise to a 1-Lipschitz retraction π from U to g_w^v defined by following the leaves of the orthogeodesic foliation to g_w^v . The image $\pi(k_{v,w})$ is then equal to the segment of g_w^v joining p_w to the point labeled by $\sigma_\lambda(X)(v, w)$, which has length $|\sigma_\lambda(X)(v, w)|$. The lemma follows. \square

Hyperbolic shearing as train track weights. Now that we have explained how to record the shapes of $X \setminus \lambda$ (Lemma 13.1) and the shears between nearby hexagons (Definition 13.2), we can package this information together to define the *geometric shear-shape cocycle* $\sigma_\lambda(X)$ of a hyperbolic structure X .

Below, we realize the shape and shear information specified above as a weight system on a standard smoothing of a geometric train track carrying λ ; this strategy allows us to specify $\sigma_\lambda(X)$ by a finite collection of information. Once we show that the weights are well-defined and satisfy the switch conditions, we then invoke Proposition 9.5 to interpret this weight system as an (axiomatic) shear-shape cocycle (see Definition 13.8). This reinterpretation in turn makes it apparent that our initial choice of train track does not matter.

Using Construction 5.6, choose a geometric train track $\tau \subset X$ that carries λ snugly and let $\tau_{\underline{\alpha}}$ be a standard smoothing of $\tau \cup \underline{\alpha}(X)$ (see Construction 9.3). Note that the components of $\tilde{X} \setminus \widetilde{\tau_{\underline{\alpha}}}$ are in bijection with the set of hexagons \mathcal{H} , and that the assumption that τ carries λ snugly ensures that if two hexagons correspond to adjacent components of $\tilde{X} \setminus \widetilde{\tau_{\underline{\alpha}}}$ then they either share an edge of $\underline{\alpha}$ or form a simple pair. We recall that two hexagons are a simple pair if the geodesic connecting their basepoints passes only through spikes of $S \setminus \lambda$.

Construction 13.5. Fix $\tau_{\underline{\alpha}} \subset X$ as above. We then associate a weight system $w(X) \in \mathbb{R}^{b(\tau_{\underline{\alpha}})}$ as follows:

- To each branch corresponding to $\alpha \in \underline{\alpha}$, assign the weight $c_\alpha = i(\mathcal{O}_\lambda(X), e_\alpha)$, where e_α is the edge of Sp dual to α .
- For each branch $b \subset \tau_{\underline{\alpha}}$ that does not correspond to an arc of $\underline{\alpha}$, choose a lift $\tilde{b} \in \widetilde{\tau_{\underline{\alpha}}}$. Let $v, w \in \mathcal{H}$ denote the vertices of $\widetilde{\text{Sp}}$ corresponding to the hexagons adjacent to \tilde{b} , and set $w(X)(b) = \sigma_\lambda(X)(v, w)$.

Lemma 13.6. Let $X, \lambda, \underline{\alpha}$ and $\tau_{\underline{\alpha}}$ be as above. Then the edge weights $w(X) \in \mathbb{R}^{b(\tau_{\underline{\alpha}})}$ given by Construction 13.5 satisfy the switch conditions.

Proof. Reference to Figure 19 will provide clarity throughout. We note that $\tau_{\underline{\alpha}}$ is generically trivalent, but may be 4-valent if there are arcs $\alpha_1, \alpha_2 \in \underline{\alpha}$ whose endpoints on λ lie on a common leaf of the orthogeodesic foliation. We give an argument only for the trivalent switches of $\tau_{\underline{\alpha}}$, and leave it to the reader to make the necessary adjustments for 4-valent switches (the statement for 4-valent switches can also be deduced by continuity).

Let s be a trivalent switch; then standing at s and looking into the spike, there are small half-branches exiting s on our right and left; call these r and ℓ , respectively. By our convention on standard smoothings, every half-branch of $\tau_{\underline{\alpha}}$ corresponding to an arc of $\underline{\alpha}$ is a right small half-branch.

If no branch of s corresponds to an arc of $\underline{\alpha}$, then the arguments appearing in [Bon96, Section 2] imply that the weights satisfy the switch conditions, because the orthogeodesic foliation is equivalent to the horocycle foliation in near s . See also [Pap91, Section 6] for a discussion more similar in spirit to ours.

Otherwise, the right small half-branch r is labeled by some $\alpha \in \underline{\alpha}$. Let b be the large half-branch incident to s . Give names also to the hexagons incident to s and their distinguished points on b or ℓ by projection; they are N, SE , and $SW \in \mathcal{H}$, and p_N, p_{SE}, p_{SW} respectively, where b and ℓ form part of the boundary of N , ℓ and r form part of the boundary of SE , and r and b form part of the boundary of SW . See Figure 19.

Now take $d = d(p_{SW}, p_{SE})$, which is equal to $w(X)(r) = c_\alpha > 0$ by definition. Define also

$$s_1 := |w(X)(b)| = d_\tau(p_{SW}, p_N) \text{ and } s_2 := |w(X)(\ell)| = d_\tau(p_N, p_{SE}).$$

Here, d_τ is understood to mean the distance between leaves of the orthogeodesic foliation near τ , measured along any leaf of λ (see Section 5.2 for an explanation of why this value is well-defined).

There are 3 kinds of configurations for the projection points p_{SW}, p_N and p_{SE} that determine the signs of $w(X)(b)$ and $w(X)(\ell)$:

- (1) The point p_N precedes both p_{SW} and p_{SE} with respect to the orientation of τ on induced by H_N , so that

$$w(X)(b) = -s_1 \text{ and } w(X)(\ell) = -s_2 \text{ with } s_2 > s_1.$$

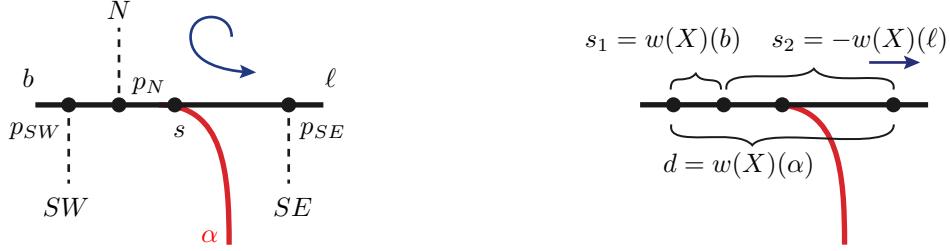


FIGURE 19. Left: A local picture of τ near s . Right: Case (3). The switch condition is satisfied because $s_1 = d - s_2$.

In this case, $d = s_2 - s_1$ and so $d - s_2 = -s_1$, which is exactly the switch condition.

- (2) Both p_{SW} and p_{SE} precede p_N , so that

$$w(X)(b) = s_1 \text{ and } w(X)(\ell) = s_2 \text{ with } s_1 > s_2.$$

This possibility gives that $d = s_1 - s_2$ and so $d + s_2 = s_1$.

- (3) The point p_{SW} precedes p_N which in turn precedes p_{SE} , so that $w(X)(b) = s_1$ and $w(X)(\ell) = -s_2$. In this case, $d = s_1 + s_2$ and so $d - s_2 = s_1$, which is again the switch condition.

Therefore, no matter the configuration of points p_N, p_{SW} , and p_{SE} , we see that the switch conditions are fulfilled at s , completing the proof of the lemma. \square

Remark 13.7. It is important to note that $w(X)$ is generally *not* the same as the weight system coming from the shear coordinates of a completion of λ (unless λ was maximal to begin with).

Invoking Proposition 9.5 and Lemma 13.1, we see that the weight system $w(X)$ defines a shear-shape cocycle with underlying arc system $\underline{A}(X)$.

Definition 13.8. The *geometric shear-shape cocycle* $(\sigma_\lambda(X), \underline{A}(X))$ of a hyperbolic metric X is the unique shear-shape cocycle for λ corresponding to the weight system $w(X)$ of Construction 13.5.

The rule that assigns to a hyperbolic structure its geometric shear-shape cocycle therefore defines a map

$$\begin{aligned} \sigma_\lambda : \mathcal{T}(S) &\rightarrow \mathcal{SH}(\lambda) \\ X &\mapsto \sigma_\lambda(X), \end{aligned}$$

which is the subject of the rest of this article.

Train track independence. We have employed the language of train tracks for convenience — the ties of a train track are a useful class of measurable arcs in the sense that they can be made transverse to λ and disjoint from $\underline{\alpha}$ (or record the weight associated to an arc of $\underline{\alpha}$). However, Construction 13.5 and Definition 13.8 *a priori* depend on the choice of geometric train track $\tau_{\underline{\alpha}}$ carrying λ .

Now that we have identified the weight system $w(X)$ with the shear-shape cocycle $\sigma_\lambda(X)$, however, we can invoke both the axiomatic and cohomological interpretations (Definitions 7.5 and 7.12) to see that the value of $\sigma_\lambda(X)$ on any arc k transverse to λ but disjoint from $\underline{\alpha}$ does not depend on the choice of geometric train track. Indeed, let k be any such arc; then by transverse invariance (axiom (SH1)) we may replace k with a concatenation of short geodesics, all of which are transverse to λ but disjoint from $\underline{\alpha}$. By additivity (axiom (SH2)), it therefore suffices to show that the value of $\sigma_\lambda(X)$ on any short geodesic disjoint from $\underline{\alpha}$ does not depend on the train track.

Lemma 13.9. Let k be a short enough geodesic segment on X that is transverse to λ . Lift k to an arc \tilde{k} on \tilde{X} and let v and w be the hexagons containing the endpoints of \tilde{k} ; then

$$\sigma_\lambda(X)(k) = \sigma_\lambda(X)(v, w)$$

where on the left $\sigma_\lambda(X)$ represents the axiomatic shear-shape cocycle and on the right $\sigma_\lambda(X)$ represents the shear between nearby hexagons (Definition 13.2). In particular, $\sigma_\lambda(X)(k)$ does not depend on the choice of train track employed in Definition 13.8.

In fact, the conclusion of this lemma holds for *all* simple pairs.

Proof. So long as k is short enough (shorter than all arcs of $\underline{\alpha}(X)$) we know that (v, w) is a simple pair. Using axiom (SH1), we may therefore isotope k through arcs transverse to λ but disjoint from $\underline{\alpha}$ to an arc k' , defined to be the concatenation of $k_{v,w}$, the geodesic connecting the points p_v and p_w on the boundary geodesics g_v^w and g_w^v , together with segments of the orthogeodesic foliation inside each hexagon H_v and H_w .

Let τ be a geometric train track snugly carrying λ defined with parameter ϵ ; then the collapse map $\pi : \mathcal{N}_\epsilon(\lambda) \rightarrow \tau$ takes k' to a train route on τ , hence on $\tau_{\underline{\alpha}}$. Orient k' (and hence also the train route $\pi(k')$) so that it travels from v to w . Let $v = u_1, u_2, \dots, u_N = w$ denote the sequence of hexagons corresponding to regions of $\tilde{X} \setminus \widetilde{\tau_{\underline{\alpha}}}$ bordering this train route, so that the regions corresponding to u_i and u_{i+1} both meet the same subsegment of $\pi(k')$. Let p_i denote their corresponding projections onto λ . Note that since $\pi(k')$ is carried on $\tau \prec \tau_{\underline{\alpha}}$, no pair of subsequent hexagons u_i and u_{i+1} lies in the same component of $\widetilde{\text{Sp}}$. This plus the construction of the train track implies that (u_i, u_{i+1}) is a simple pair, and we can measure the shear $\sigma_\lambda(X)(u_i, u_{i+1})$ (up to sign) as the distance along the train track between $\pi(p_i)$ and $\pi(p_{i+1})$.

Now given $\tau_{\underline{\alpha}}$ carrying λ , we observe that k' also determines a (pair of) relative cycle(s) in the corresponding (orientation cover of the) ϵ -neighborhood of $\lambda_{\underline{\alpha}}$. The value $\sigma_\lambda(X)(k) = \sigma_\lambda(X)(k')$ is then equal to the value of the cohomological shear-shape cocycle evaluated on either of the oriented lifts \hat{k}' of k' which cross the lift of λ with positive local orientation. We may therefore express

$$[\hat{k}'] = [t_1] - [t_2] + [t_3] - \dots \pm [t_{N-1}]$$

where t_i is a (lift of a) tie corresponding to the branch of the train track connecting the regions corresponding to u_i and u_{i+1} , lifted to the orientation cover to have positive intersection with λ . See Figure 20.

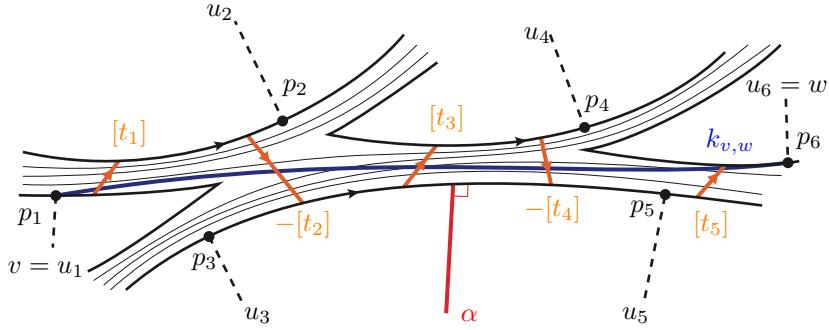


FIGURE 20. Measuring the shear of a small arc using a geometric train track. By isotoping k to a proper arc in the geometric train track neighborhood and then expressing its relative homology class as a sum of the branches, we can compute its shear as the alternating sum of shears between adjacent hexagons.

But now by construction, we know that $\sigma_\lambda(X)$ evaluated on $[t_i]$ is just the shear $\sigma_\lambda(X)(u_i, u_{i+1})$. In turn, this shear is equal to the signed distance along the train track between $\pi(p_i)$ and $\pi(p_{i+1})$ (where the sign is determined by the local orientation of λ). Combining this with the expression for $[\hat{k}']$ above, we see that $\sigma_\lambda(X)(k)$ is exactly equal to the signed distance along the train track between $\pi(p_1)$ and $\pi(p_N)$, which is the shear $\sigma_\lambda(X)(v, w)$. \square

We note that in the proof above, the cohomological interpretation of shear-shape cocycles provides a convenient workaround for the obstacle that the train route with dual transversals t_1, \dots, t_{N-1} is not in general isotopic to k through arcs transverse to λ . Regardless, the relative homology class defined by $k' \cap \mathcal{N}_\epsilon(\lambda)$ is homologous to a linear combination of $\{t_i\}$ in the orientation cover of $\mathcal{N}_\epsilon(\lambda)$.

Remark 13.10. The lemma above can also be proved by splitting any two geometric train tracks to a common subtrack [PH92, Theorem 2.3.1]. Each splitting sequence can then be realized in the orthogeodesically-foliated neighborhood $\mathcal{N}_\epsilon(\lambda) \subset X$ by cutting along compact paths in the spine associated to a spike, as in [ZB04, Section 3]. Splits induce maps on weight spaces, and so Lemma 13.9 is essentially equivalent to the statement that Construction 13.5 is compatible with splitting and collapsing. See also [Bon97b, Lemma 6].

The cocycle as a map on pairs. It will be convenient to repackage the data provided by $\sigma_\lambda(X)$ in yet another form, which also explains our choice of notation in Definition 13.2.

If $v, w \in \mathcal{H}$ can be joined by a Lipschitz continuous segment $k_{v,w}$ which is transverse to λ , disjoint from α , and meets no leaf of $\tilde{\lambda}$ twice, then we say that (v, w) is a *transverse pair* and that $k_{v,w}$ is a *transversal*. If (v, w) is a transverse pair, we say that r is *between* v and w if there is a transversal $k_{v,w}$ that decomposes as a concatenation of transversals $k_{v,w} = k_{v,r} \cdot k_{r,w}$. Finally, we define

$$\sigma_\lambda(X)(v, w) := \sigma_\lambda(X)(k_{v,w})$$

and declare that $\sigma_\lambda(X)(v, v) = 0$. Observe that if (v, w) are a simple pair then this agrees with our definition of the shear between nearby hexagons (Definition 13.2).

Lemma 13.11. The shear-shape cocycle $\sigma_\lambda(X)$ defines a map on transverse pairs that satisfies

- (1) (π_1 -invariance) for each $\gamma \in \pi_1(X)$, we have $\sigma_\lambda(X)(\gamma v, \gamma w) = \sigma_\lambda(X)(v, w)$.
- (2) (finite additivity) If (v, w) is a transverse pair and r is between v and w , then

$$\sigma_\lambda(X)(v, w) = \sigma_\lambda(X)(v, r) + \sigma_\lambda(X)(r, w).$$

- (3) (symmetry) $\sigma_\lambda(X)(v, w) = \sigma_\lambda(X)(w, v)$.

The proof of this lemma is simply a consequence of unpacking the definitions and showing that two different choices of transversals give the same shear values; the latter statement is just a repeated application of Axiom (SH3).

13.3. Injectivity and positivity. We now record some initial structural properties of the map σ_λ defined above. In particular, we demonstrate that σ_λ is injective and interacts coherently with the orthogeodesic foliation map \mathcal{O}_λ and the shear-shape coordinatization I_λ of transverse foliations.

Observe that injectivity of σ_λ is equivalent to the statement that if two hyperbolic structures have the same complementary subsurfaces and same gluing data along λ , then they must be isometric. As the horocyclic and orthogeodesic foliations are equivalent in spikes of complementary subsurfaces, the proofs of [Bon96, Lemma 11 and Theorem 12] may be invoked *mutatis mutandis*. We outline this argument below for the convenience of the reader, and direct them to [Bon96] for a more thorough discussion of the estimates involved. We remark that this strategy also appears in the proof of Proposition 15.12, where we use it to piece together Lipschitz-optimal homeomorphisms along λ .

Proposition 13.12. The map $\sigma_\lambda : \mathcal{T}(S) \rightarrow \mathcal{SH}(\lambda)$ is injective.

Sketch of proof. Fix homeomorphisms $(\tilde{S}, \tilde{\lambda})$ with $(\tilde{X}_i, \tilde{\lambda})$ that lift the markings $S \rightarrow X_i$ and so that each component $\tilde{\Sigma} \subset \tilde{S} \setminus \tilde{\lambda}$ maps homeomorphically to a component $\tilde{Y}_i \subset \tilde{X}_i \setminus \tilde{\lambda}$ for $i = 1, 2$.

Suppose that $\sigma_\lambda(X_1) = \sigma_\lambda(X_2)$; then in particular $A(X_1) = A(X_2)$ and so by Theorem 6.4, the complementary subsurfaces $\overline{X_1 \setminus \lambda}$ and $\overline{X_2 \setminus \lambda}$ are isometric. Therefore, for a given component $\Sigma \subset S \setminus \lambda$, we can find an $\pi_1(\Sigma)$ equivariant isometry $\varphi_\Sigma : \tilde{Y}_1 \rightarrow \tilde{Y}_2$. Define $\varphi : \tilde{X}_1 \setminus \lambda \rightarrow \tilde{X}_2 \setminus \lambda$ to be the union of these maps on each complementary component; by construction, φ is an isometry.

We need to show that φ extends to a $\pi_1(S)$ -equivariant isometry $\varphi : \tilde{X}_1 \rightarrow \tilde{X}_2$. To prove this, we apply the arguments of [Bon96, Lemma 11], which we summarize presently. The first step is to construct a locally Lipschitz continuous extension of φ ; this step employs the length bound of Lemma 13.4 and some elementary hyperbolic geometry, and the arguments of the first ten paragraphs of [Bon96, Lemma 11] may be applied verbatim.

As in Bonahon's original proof, we now show that φ is actually 1-Lipschitz, given that it is locally Lipschitz. We first show that φ does not increase the length of leaves of the orthogeodesic foliation.

Given any segment ℓ of a leaf of the orthogeodesic foliation $\mathcal{O}_\lambda(X_1)$, the length of ℓ restricted to any hexagon H_u where $u \in \mathcal{H}$ is completely determined by the isometry type of H_u and the distance along $\tilde{\lambda}$ from $p_u \in \partial_\lambda H_u$. As $\sigma_\lambda(X_1)$ determines the shape of $X_1 \setminus \lambda$, we can recover this information and hence determine the length of $\ell \cap H_u$ just from the data of $\sigma_\lambda(X_1)$.

From $\sigma_\lambda(X_2) = \sigma_\lambda(X_1)$, we deduce that the length of ℓ in any hexagon of \tilde{X}_1 is equal to the length of $\varphi(\ell)$ in the corresponding hexagon of \tilde{X}_2 . Moreover, since φ is locally Lipschitz, the 1-dimensional Lebesgue measure of $\varphi(\ell) \cap \varphi(\tilde{\lambda})$ is at most the 1-dimensional Lebesgue measure of $\ell \cap \lambda$. By a now classical fact the

latter is zero [BS85], hence so is the former. Therefore, the length of ℓ in X_1 is equal to the length of ℓ in X_2 .

Now there is a path joining any two points in \tilde{X}_1 built from geodesic segments and segments of leaves of the orthogeodesic foliation. The argument above shows that φ preserves the lengths of such paths, so φ is globally 1-Lipschitz. The construction is completely symmetric, so φ^{-1} is 1-Lipschitz as well. Now every 1-Lipschitz homeomorphism between metric spaces with 1-Lipschitz inverse is necessarily an isometry, and equivariance of φ is immediate from the construction. Therefore X_1 and X_2 must be isometric. \square

The diagram commutes. We have now developed sufficient technology to prove that the geometric shear-shape cocycle of a hyperbolic metric is the same as the shear-shape cocycle associated to its orthogeodesic foliation. In other words, Diagram (2) commutes. Compare with [Mir08, Proposition 6.1].

Theorem 13.13. For all $\lambda \in \mathcal{ML}$ and all $X \in \mathcal{T}(S)$ we have $\sigma_\lambda(X) = I_\lambda \circ \mathcal{O}_\lambda(X)$.

Proof. Fix a standard smoothing $\tau_{\underline{\alpha}}$ of a geometric train track τ for λ on X . Our approach is to compute both $\sigma_\lambda(X)(b)$ and $I_\lambda \circ \mathcal{O}_\lambda(X)(b)$ for each branch b of $\tau_{\underline{\alpha}}$. These numbers will coincide, so by Proposition 9.5, $\sigma_\lambda(X) = I_\lambda \circ \mathcal{O}_\lambda(X)$.

Let $T_X \subset X$ be the piecewise geodesic triangulation of X whose vertices are the vertices of \mathbf{Sp} , so that there is an edge between $v, w \in \mathbf{Sp}$ if the corresponding regions of $X \setminus \tau_{\underline{\alpha}}$ share a branch. This recipe generically yields a triangulation, but may have quadrilaterals in the case that two points of $\underline{\alpha}(X) \cap \lambda$ lie on the same leaf of $\mathcal{O}_\lambda(X) \cap \mathcal{N}_\epsilon(\lambda)$. In this case, we may either choose a smaller initial neighborhood to define our geometric train track so that this does not occur, or these points correspond to arcs that meet an isolated leaf of λ on either side; in the latter case, choose either diagonal that crosses the quadrilateral to include into T_X . Observe that each edge of T_X is either transverse to $\mathcal{O}_\lambda(X)$ or a segment of a leaf (on the off chance that two adjacent regions have exactly 0 shear between them).

Let $q = q(\mathcal{O}_\lambda(X), \lambda)$, and recall that Proposition 5.10 provides a homotopy equivalence $\mathcal{D}_\lambda : X \rightarrow q$ in the correct homotopy class satisfying $\mathcal{D}_{\lambda*} \mathcal{O}_\lambda(X) = V(q)$ and $\mathcal{D}_{\lambda*} \lambda = H(q)$ both leafwise and measurably. Furthermore, \mathcal{D}_λ maps T_X to a (topological) triangulation of q with vertices at its zeros. It therefore remains to show that $\sigma_\lambda(X)$ evaluated on a branch of $\tau_{\underline{\alpha}}$ is the same as $I_\lambda(q)$ evaluated on the dual edge of this triangulation.

Now by definition, $\underline{A}(X) = \underline{A}(q)$, so consider a branch b of $\tau_{\underline{\alpha}}$ not corresponding to an arc of the arc system. Dual to b there is an edge e of the triangulation $\mathcal{D}_\lambda(T_X)$ which is transverse to the orthogeodesic foliation $\mathcal{O}_\lambda(X)$ on q (since T_X was transverse to $\mathcal{O}_\lambda(X)$ on X). Up to sign, the value of $I_\lambda(q)$ on b is the magnitude of the real part of the period of e , which is just the geometric intersection number $i(\mathcal{O}_\lambda(X), e)$ by transversality.

On the other hand, we have that $\sigma_\lambda(X)(b)$ is equal to the shear between the two hexagons on either side of b . This in turn is equal to the geometric intersection number $i(\mathcal{O}_\lambda(X), k_{v,w})$ up to sign, where $k_{v,w}$ is the geodesic connecting the vertices p_v and p_w of $\lambda \cap \underline{\alpha}(X)$. Since \mathcal{D}_λ takes $k_{v,w}$ to an arc transversely isotopic to e , we have that $|\sigma_\lambda(X)(b)| = |I_\lambda(q)(b)|$.

Finally, to show that the signs are equal, fix matching orientations on $k_{v,w}$ and e . These induce a local orientations on the leaves of λ so that the algebraic intersection of λ with $k_{v,w}$, respectively e , is positive. In turn, this induces a local orientation on the leaves of $\mathcal{O}_\lambda(X)$ near $k_{v,w}$, respectively e , and our sign conventions are equivalent to stipulating that the sign is positive if $k_{v,w}$, respectively e , crosses $\mathcal{O}_\lambda(X)$ from left to right and negative if it crosses from right to left (compare [Mir08, §5.2]). In particular, the signs agree and so $\sigma_\lambda(X)(b) = I_\lambda(q)(b)$ for all branches b , completing the proof of the theorem. \square

Corollary 13.14. For all $\mu \in \Delta(\lambda)$, we have an equality

$$\omega_{\mathcal{SH}}(\sigma_\lambda(X), \mu) = i(\mathcal{O}_\lambda(X), \mu) = \ell_X(\mu) > 0.$$

In particular, $\sigma_\lambda(\mathcal{T}(S)) \subseteq \mathcal{SH}^+(\lambda)$.

Proof. The first equality is a direct consequence of Theorem 13.13 and Proposition 10.12. The second equality was proved in Lemma 5.7. \square

14. SHAPE-SHIFTING COCYCLES

In the previous section, we explained how to associate to each hyperbolic structure X a shear-shape cocycle $\sigma_\lambda(X)$. In this one, we explain how to upgrade a small deformation \mathfrak{s} of the cocycle into a deformation of the hyperbolic structure; this is eventually used to prove that $\sigma_\lambda : \mathcal{T}(S) \rightarrow \mathcal{SH}^+(\lambda)$ is open (Theorem 15.1 below). The main issue that we need to overcome is that we must simultaneously change the geometry of the non-rigid components of $X \setminus \lambda$ while shearing these subsurfaces along one another.

The goal of this section is therefore to build, for every small enough deformation \mathfrak{s} of $\sigma_\lambda(X)$, a $\pi_1(S)$ -equivariant *shape-shifting cocycle* that records how to adjust the relative position of geodesics of λ :

$$\varphi_{\mathfrak{s}} : \partial_\lambda \mathcal{H} \times \partial_\lambda \mathcal{H} \rightarrow \text{Isom}^+ \tilde{X}$$

where $\partial_\lambda \mathcal{H} := \{(h_v, p_v) \subset \partial_\lambda H_v : v \in \mathcal{H}\}$ is the set of boundary geodesics of $\tilde{\lambda}$ equipped with basepoints obtained from projections of the vertices of $\tilde{\mathcal{S}}\mathcal{P}$. See Proposition 14.26.

In Section 15.1 below, we explain how to modify the developing map $\tilde{X} \rightarrow \mathbb{H}^2$ according to $\varphi_{\mathfrak{s}}$, resulting in a new (equivariant) hyperbolic structure $X_{\mathfrak{s}}$ with geometric shear-shape cocycle $\sigma_\lambda(X) + \mathfrak{s}$ (Lemma 15.6). By fixing a pointed geodesic $(h_v, p_v) \in \partial_\lambda \mathcal{H}$ we identify $\text{Isom}^+(\tilde{X})$ with $T^1 \tilde{X}$, so that the projection of $\{\varphi_{\mathfrak{s}}((h_v, p_v), (h_w, p_w)) \mid (h_w, p_w) \in \partial_\lambda \mathcal{H}\}$ to \tilde{X} is then be the geodesic realization of $\tilde{\lambda}$ in the new metric $\tilde{X}_{\mathfrak{s}}$.

When the deformation \mathfrak{s} preserves $\underline{A}(X)$, the cocycle $\varphi_{\mathfrak{s}}$ corresponds to a cataclysm map: the complementary components of $\tilde{X} \setminus \tilde{\lambda}$ are sheared along the leaves of $\tilde{\lambda}$ and map isometrically into the deformed surface $X_{\mathfrak{s}}$. When \mathfrak{s} alters $\underline{A}(X)$, we must shear the complementary subsurfaces while also simultaneously changing their shape, introducing complications not present in Bonahon and Thurston's original considerations.

Deforming the cocycle. We first make explicit what we mean by a deformation of a shear-shape cocycle; we quantify what we mean by “small” in Section 14.2.

Observe that if σ and σ' in $\mathcal{SH}^+(\lambda)$ are close, then by Proposition 8.5 we know that their underlying weighted arc systems \underline{A} and \underline{A}' are close in $\mathcal{B}(S \setminus \lambda)$. In particular, the corresponding unweighted arc systems $\underline{\alpha}$ and $\underline{\alpha}'$ must both live in some common top-dimensional cell of $\mathcal{B}(S \setminus \lambda)$, i.e., must both be contained in some common maximal arc system $\underline{\beta}$. Let τ be some snug train track for λ and let $\tau_{\underline{\beta}}$ be a standard smoothing of $\tau \cup \underline{\beta}$. By Proposition 9.5, we may then identify σ and σ' as weight systems on $\tau_{\underline{\beta}}$; the difference $\sigma - \sigma' \in W(\tau_{\underline{\beta}})$ is then a deformation of σ .

In general, if $(\sigma, \underline{A}) \in \mathcal{SH}^+(\lambda)$ and $\underline{\beta}$ is any maximal arc system containing the support of \underline{A} , then the deformations we consider in this section are those $\mathfrak{s} \in W(\tau_{\underline{\beta}})$ such that $\sigma + \mathfrak{s} \in W(\tau_{\underline{\beta}})$ corresponds to a positive shear-shape cocycle. Passing between equivalent definitions of shear-shape cocycles, we see that we may also think of \mathfrak{s} as a “shear-shape cocycle with negative arc weights.” The underlying weighted arc system of any deformation \mathfrak{s} will be denoted by \mathfrak{a} ; while its coefficients are not necessarily positive, they will satisfy the zero total residue condition of (13) by construction.

By Theorem 6.4, the arc system $\underline{A} + \mathfrak{a}$ gives each component of $S \setminus \lambda$ a new complete hyperbolic metric Y with (non-compact) totally geodesic boundary. Since the supports of \underline{A} and $\underline{A} + \mathfrak{a}$ are both contained inside of some common maximal $\underline{\beta}$, one may set up a correspondence between the complementary components of $X \setminus \lambda_{\underline{\alpha}}$ with the components of $Y \setminus \text{supp}(\underline{A} + \mathfrak{a})$ (adding in weight 0 edges as necessary).

A blueprint. To help guide the reader through this rather intricate construction, we include here a top-level overview of the necessary steps, together with an outline of the section. Briefly, our strategy is to explicitly define $\varphi_{\mathfrak{s}}$ on two types of pairs of pointed geodesics: the “simple pairs” between which the orthogeodesic foliation is comparable to the horocyclic, and the pairs which live in the boundary of a common subsurface. Piecing together these basic deformations then allows us to define $\varphi_{\mathfrak{s}}$ on arbitrary pairs of pointed geodesics.

Our construction of $\varphi_{\mathfrak{s}}$ for simple pairs parallels Bonahon's construction of shear maps [Bon96, Section 5], and as such requires a detailed analysis of the geometry of the spikes of $\tilde{X} \setminus \tilde{\lambda}$. We therefore devote Section 14.1 to recording a number of useful notions and estimates from [Bon96]. In this section, we also introduce the “injectivity radius of X along λ ,” which measures the length of the shortest curve carried on a maximal snug train track for λ and plays a crucial role in our convergence estimates.

After these preliminary considerations, we turn in Section 14.2 to the actual construction of $\varphi_{\mathfrak{s}}$ on simple pairs. As in [Bon96], the map is defined by adjusting the lengths of countably many horocyclic arcs in an appropriate neighborhood of λ , compensating for changing shears between hexagons. Unlike in [Bon96], we must also adjust the arcs to account for the changing shapes of each of the spikes (as we are deforming

the complementary subsurfaces). Convergence of the resulting infinite product of parabolic transformations is delicate; our approach follows [Bon96, Section 5] with influence from the more geometric approach of [Thu98]. An accessible treatment of Thurston's construction of “cataclysm coordinates” can be found in [PT07, Section 3.5].

We then turn in Sections 14.3 and 14.4 to defining φ_s on pairs of geodesics in the boundary of the same hexagon or the same complementary subsurface, respectively. It is here that our work significantly differs from that of Bonahon and Thurston. In these sections we also develop the idea of “sliding” a deformed complementary subsurface along the original; this viewpoint allows us to easily demonstrate a number of otherwise nontrivial relations between Möbius transformations (see Propositions 14.18, 14.19, and 14.24).

Finally, in Section 14.5 we build the shape-shifting cocycle φ_s from these pieces; the cocycle relation (Proposition 14.26) then follows from the cocycle relations for pieces and the separation properties of $\tilde{\lambda}$.

Note. We remark that throughout this section and the next, we consider isometries via their action on a pointed geodesic, and compositions should be read from right to left.

14.1. Geometric control in the spikes. We first record some useful definitions and associated geometric estimates. These estimates play a crucial role in establishing convergence of the infinite products appearing in Section 14.2 below. Many of our definitions follow Bonahon's, but in order to contend with the fact that the complementary subsurfaces of λ are not always isometric, we must relate certain constants to the geometry of λ on X (see Lemma 14.5 in particular).

Our discussion will take place with certain data fixed. Choose a hyperbolic surface $X \in \mathcal{T}(S)$ and a measured lamination $\lambda \in \mathcal{ML}(S)$. Let $\epsilon > 0$ be small enough so that an ϵ -geometric train track τ on X carries λ snugly. The standard smoothing $\tau_{\underline{\alpha}}$ for the arc system $\underline{\alpha} = \underline{\alpha}(X)$ provides us with a vector space $W(\tau_{\underline{\alpha}})$ that models $\mathcal{SH}(\lambda; \underline{\alpha})$. With $\tau_{\underline{\alpha}}$ fixed, we endow the vector space of weights on branches of $\tau_{\underline{\alpha}}$ with the sup norm $\|\cdot\|_{\tau_{\underline{\alpha}}}$, and restrict this norm to the weight space $W(\tau_{\underline{\alpha}})$.

Let k_b be an oriented geodesic transverse to a branch $b \in \tau$ that also avoids $\underline{\alpha}$. Following Bonahon, we define the *divergence radius* or *depth* $r_b(d) \in \mathbb{Z}_{>0}$ of a component d of $k_b \setminus \lambda$ to be “how long the leaves of λ incident to d track each other,” as viewed by τ .

More precisely, lift everything to the universal cover \tilde{X} . By convention, set $r_b(d) = 1$ if d contains one of the endpoints of k_b . Otherwise, d is contained in a spike of H_v for some $v \in \mathcal{H}$, i.e., d connects a pair of asymptotic geodesics g_d^- and g_d^+ . The divergence radius $r_b(d)$ is then the largest integer $r \geq 1$ such that $\pi(g_d^+)$ and $\pi(g_d^-)$ successively cross the same sequence of branches

$$b_{-r+1}, b_{-r+2}, \dots, b_0, \dots, b_{r-2}, b_{r-1}$$

of $\tilde{\tau}$, where b_0 is the lift of b meeting \tilde{k}_b and $\pi : \mathcal{N}_\epsilon(\tilde{\lambda}) \rightarrow \tilde{\tau}$ is the collapse map. By equivariance, $r_b(d)$ is clearly independent of the choice of lift \tilde{k}_b of k_b .

Remark 14.1. After projecting back down to $\tau \subset X$, either $b_{-r+1} \cdot \dots \cdot b_0$ or $b_0 \cdot \dots \cdot b_{r-1}$ defines a train route γ_d in τ that starts at b and terminates by “opening up” into the projection of H_v in X . That is, the geodesics g_d^+ and g_d^- diverge from each other (at scale ϵ) at the terminus of γ_d .

Now there are boundedly many spikes of $X \setminus \lambda$, and for each $r \geq 1$ each spike may contain at most 1 component $d \subset k_b \setminus \lambda$ with depth exactly r . This gives us the following bound:

Lemma 14.2 (Lemma 4 of [Bon96] and Lemma 5 of [SB01]). For any branch b of τ and any transversal k_b , the number of components d of $k_b \setminus \lambda$ with $r_b(d) = r$ is at most $6|\chi(S)|$.

The train track interpretation of the depth of a segment also allows us to bound the value of a shear-shape cocycle s in terms of its weights on a snug train track and the depth of its endpoints.

More specifically, for each component d of $k_b \setminus \lambda$, let k_b^d be the subarc of k_b joining the initial point of k_b to any point of d . Then for any combinatorial deformation s and b a branch of $\tau_{\underline{\alpha}}$, there is an explicit formula for $s(k_b^d)$ as a linear function of the weights of s on $\tau_{\underline{\alpha}}$ with at most $r_b(d)$ terms [Bon97b, Lemma 6]. Conceptually, this formula arises by splitting $\tau_{\underline{\alpha}}$ open along the spike s containing d , until d is “visible” in some new track $\tau'_{\underline{\alpha}}$ carried by $\tau_{\underline{\alpha}}$ (see also the proof of Lemma 13.9).

The exact expression for $s(k_b^d)$ will not be important for us; instead, we record the following estimate, which follows by considering the growth of edge weights upon splitting.

Lemma 14.3 (Lemma 6 of [Bon96] and Lemma 6 of [SB01]). Let k_b be a transversal of a branch b . Then

$$|\mathfrak{s}(k_b^d)| \leq \|\mathfrak{s}\|_{\tau_{\underline{\alpha}}} r_b(d)$$

for every $\mathfrak{s} \in \mathcal{SH}(\lambda; \underline{\alpha})$ and every component d of $k_b \setminus \lambda$.

We remark that our definitions of $\|\cdot\|_{\tau_{\underline{\alpha}}}$ and $r_b(\cdot)$ make the bound given in Lemma 14.3 hold without a topological multiplicative factor, as in [Bon96].

Geometric estimates on depth. The depth of a component d of $k_b \setminus \lambda$ is proportional to the distance from a lift \tilde{d} to the vertex $u \in \mathcal{H}$ inside of the corresponding spike. The constant of proportionality in turn depends on how quickly the spike of H_u containing \tilde{d} returns to k_b on X ; we now identify a quantity that will allow us to estimate this constant.

Let k be any geodesic arc transverse to λ such that each lift \tilde{k} to \tilde{X} bounds a spike in every hexagon that it crosses; equivalently, the endpoints of \tilde{k} lie in a simple pair of hexagons. As in Section 13.2, it suffices for k to be shorter than the shortest arc of $\underline{\alpha}(X)$. Now for each leaf g of $\tilde{\lambda}$, there is a bound $R_k(g) > 0$ for the distance in g between intersections of g with different lifts \tilde{k}_1 and \tilde{k}_2 of k . Indeed, any two lifts of k meeting g differ by a deck transformation $\gamma \in \pi_1(X)$ determined by a path in X that traces along the projection of a segment in g and then closes up along k .

We then define the *injectivity radius of X along λ* to be

$$\text{inj}_\lambda(X) := \inf_{k \pitchfork \lambda} \inf_{g \subset \tilde{\lambda}} R_k(g)$$

where the infimum is taken over all transverse arcs k whose endpoints lie in a simple pair of hexagons.

Equivalently, the injectivity radius of λ may also be computed by taking a ϵ so that the geometric train track τ_{\max} built from $\mathcal{N}_\epsilon(\lambda)$ is snug and so that for all $\epsilon' > \epsilon$, the train track built from $\mathcal{N}_\epsilon(\lambda)$ is the same (not just equivalent) to τ_{\max} , as follows.¹⁹

For each branch of τ_{\max} , choose a tie t_b (that is, a leaf of the orthogeodesic (or horocyclic) foliation restricted to $\mathcal{N}_\epsilon(\lambda)$ that is transverse to b). The injectivity radius along λ is then equal to the infimum of the recurrence times of λ to any t_b . Using the “length along a geometric train track” function $\ell_{\tau_{\max}}$ defined in Section 5.2, we may therefore write

$$(29) \quad \text{inj}_\lambda(X) = \inf_{\gamma \prec \tau_{\max}} \ell_{\tau_{\max}}(\gamma)$$

where the infimum is taken over all simple closed curves γ carried on the train track τ_{\max} .

Remark 14.4. The length of the hyperbolic systole of X is clearly a lower bound for $\text{inj}_\lambda(X)$, which is therefore positive. However, $\text{inj}_\lambda(X)$ can be much larger than the length of the systole.

For example, if λ does not fill the surface then there can be a disjoint curve of arbitrarily small length. In addition, X may have a very short curve γ transverse to λ , and if λ does not twist around γ , then $\text{inj}_\lambda(X)$ is necessarily very large.

We can now relate the geometry of small arcs to their depth and injectivity radius along λ .

Lemma 14.5 (Lemmas 3 and 5 of [Bon96] and Lemma 4 of [SB01]). Given a branch b of a geometric train track τ constructed from λ on X and a short transversal k_b , there exists $B > 0$ such that the following holds. For every component d of $k_b \setminus \lambda$ with depth $r_b(d)$,

$$\ell_X(d) \leq B e^{-D_\lambda(X)r_b(d)},$$

where $D_\lambda(X) = \text{inj}_\lambda(X)/|\chi(S)|$.

Proof. The idea is the same as in the reference, but our constants are different. Small geodesic arcs meeting a spike s of a hexagon H_v transversely and far away from the vertex v look like horocycles, which have length that decays exponentially in distance from v . Therefore, we just need to give a lower bound for the distance between d and $v \in H_v$ along the spike s in terms of $\text{inj}_\lambda(X)$ and the topological complexity of S .

Consider the train path γ_d starting at b that defines $r_b(d)$. By definition, γ_d traverses exactly $r_b(d)$ branches of τ (counted with multiplicity). Now γ_d decomposes as a concatenation of maximal sub-train paths with embedded interiors, each forming a *simple loop* in τ .²⁰

¹⁹Observe that any ϵ sufficiently close to the supremum of ϵ for which $\mathcal{N}_\epsilon(\lambda)$ is snug satisfies these conditions.

²⁰A simple loop on a train track is a carried curve which traverses each branch at most once.

The depth $r_b(d)$ is thus bounded above by the number of consecutive simple loops in γ_d times the size of the longest simple loop in τ . The size of a simple loop in τ is in turn bounded above by the number of branches of τ , which is at most $9|\chi(S)|$. Finally, since each simple loop in γ_d is carried on $\tau \prec \tau_{\max}$ it must have length at least $\text{inj}_\lambda(X)$ by (29).

Putting the above estimates together, we see that the distance between v and d in H_v is at least

$$\text{inj}_\lambda(X) \cdot \#\{\text{simple loops in } \gamma_d\} \geq \frac{\text{inj}_\lambda(X)r_b(d)}{\text{size of the longest simple loop in } \tau} \geq \frac{\text{inj}_\lambda(X)}{9|\chi(S)|}r_b(d),$$

and the lemma follows. \square

14.2. Shape-shifting in the spikes. Our discussion now begins to diverge from [Bon96]. While pairs of asymptotic geodesics are all isometric, the spikes of $X \setminus \lambda$ come with extra decoration, namely, a choice of horocycle at each cusp (equivalently, basepoints which lie on a common leaf of the orthogeodesic foliation). In this section, we explain how to use these decorations to define the shape-shifting cocycle φ_s on pairs of basepointed geodesics coming from simple pairs of hexagons.

We remind the reader that X , λ , and $\tau_{\underline{\alpha}}$ are fixed so that geometric objects like geodesic segments, hexagons, arcs of $\underline{\alpha}(X)$, etc. are understood to live in and be realized (ortho)geodesically on X . Throughout this section we will fix $\underline{A} = \underline{A}(X)$ and use it to denote both a weighted arc system as well as the induced metric on $S \setminus \lambda$. Finally, we recall that \mathfrak{s} is a combinatorial deformation of $\sigma_\lambda(X)$ which changes \underline{A} by \mathfrak{a} ; we will refer to the deformed hyperbolic structure on $S \setminus \lambda$ by $\underline{A} + \mathfrak{a}$ and its hexagonal pieces by G_u for $u \in \mathcal{H}$.

Shapes of spikes. The group $\text{PSL}_2(\mathbb{R})$ acts transitively on pairs of asymptotic geodesics but, having done so, cannot further act on the family of horocycles based at the spike. To measure this failure, we associate below a geometric parameter which records the placement of basepoints in each spike.

Suppose that $u \in \mathcal{H}$ is a hexagon of $\tilde{X} \setminus \widetilde{\lambda_{\underline{\alpha}}}$ and s is a spike of H_u , that is, a pair of asymptotic geodesics g and g' . Both g and g' come with basepoints p and p' obtained by projecting u to these geodesics. We then associate to s the number $h_{\underline{A}}(s)$ which measures the length of either of the orthogeodesic leaves connecting u to p or p' :

$$h_{\underline{A}}(s) := d(p, u) = d(p', u).$$

Our notation reflects the fact that this function clearly depends only on the geometry of $X \setminus \lambda$ and not the shearing along λ . The reader familiar with the literature will observe that this parameter is essentially an orthogeodesic version of the “sharpness functions” appearing in [Thu98].

In order to measure the difference in sharpness functions between the realizations of s in \underline{A} and in the deformed metric $\underline{A} + \mathfrak{a}$, we superimpose the hexagons H_u and G_u and measure the distance between their boundary basepoints.

More concretely, choose an arbitrary orientation \vec{s} of the spike s and fix realizations of both H_u and G_u inside of \mathbb{H}^2 . As $\text{PSL}_2(\mathbb{R})$ acts simply transitively on triples in $\partial\mathbb{H}^2$, there is a unique isometry that takes the realization of s in G_u to its realization in H_u . The vertex u of Sp is realized in both H_u and G_u ; let p and q denote the projections of these points to one of the boundary geodesics g of s . See Figure 21.

Lemma 14.6. With all notation as above, the signed distance from q to p along g is

$$(30) \quad f_{X, \mathfrak{s}}(\vec{s}) := \varepsilon \log \left(\frac{\tanh h_{\underline{A}}(s)}{\tanh h_{\underline{A} + \mathfrak{a}}(s)} \right) \in \mathbb{R},$$

where $\varepsilon = +1$ if \vec{s} is oriented towards the shared ideal endpoint, and $\varepsilon = -1$ otherwise.

The parameter $f_{X, \mathfrak{s}}(\vec{s})$ plays a crucial role below in our definition of the shape-shifting map on spikes. In our convergence estimates, we will also need to consider the parameter

$$(31) \quad \|\mathfrak{s}\|_{\vec{s}} := \max_s |f_{X, \mathfrak{s}}(\vec{s})| < \infty$$

which quantifies the maximum distance that the deformation \mathfrak{s} moves a basepoint in a spike.

Proof. We compute in the upper half plane; up to isometry, we may assume that s is bound by the imaginary axis $V = i\mathbb{R}_{>0}$ and its translate $2 + V$; the spine of the orthogeodesic foliation in this spike is a subsegment of the vertical line $1 + V$. With this choice fixed, the projections p and q of u to V may be identified with ie^a and ie^b for some a and b , respectively. Without loss of generality, we may also assume that V is

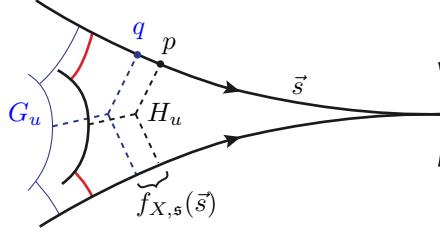


FIGURE 21. Superimposing hexagons to measure the difference in the shapes of their spikes.

oriented upwards (towards ∞); the opposite choice of orientations simply reverses all signs at the end of the computation.

Now for $t \geq 0$, the path $t \mapsto \tanh t + i \operatorname{sech} t$ is the unit speed parametrization of the orthogeodesic emanating from V at i . Observe that the isometry $z \mapsto e^a z$ stabilizes V and takes this segment to an orthogeodesic segment emanating from $ie^a = p$ which is distance a from i . Since the orthogeodesic segment through p meets the spine $1 + V$ after traveling distance $h_{\underline{A}}(s)$ (by definition), this implies that

$$e^a = \tanh h_{\underline{A}}(s).$$

Similarly, we have that $e^b = \tanh h_{\underline{A}+\mathfrak{a}}(s)$. Together, these imply that

$$\frac{\tanh h_{\underline{A}}(s_u)}{\tanh h_{\underline{A}+\mathfrak{a}}(s_u)} = e^{a-b}.$$

Taking logarithms, we see that $a - b$ is the signed distance from q to p along V , as claimed. \square

Remark 14.7. Note that by Theorem 6.4, the parameter $f_{X,s}(\vec{s})$ varies analytically in \mathfrak{a} (hence \mathfrak{s}).

Orientation conventions. We now specialize to the case where (v, w) is a simple pair of hexagons with associated oriented geodesic $k_{v,w}$ running between p_v^w on g_v^w (the projection of v to the boundary leaf of $\partial_\lambda H_v$ closest to w) and p_w^v on g_w^v .

Each leaf $g \subset \tilde{\lambda}$ crossed by $k_{v,w}$ inherits an orientation by declaring that *turning right* onto g while traveling from v to w along $k_{v,w}$ is the positive direction. We remark that if $k_{v,w}$ crosses a hexagon H_u , then the induced orientation of g_u^w , the geodesic in $\partial_\lambda H_u$ closest to w , is the *opposite* of the orientation of g_u^w induced as a part of the boundary of H_u . On the other hand, the two orientations on g_u^v induced by $k_{v,w}$ and coming from H_u agree. This is an artifact of our sign convention for measuring shears; see Remark 13.3.

If g is a complete oriented geodesic in the hyperbolic plane and $t \in \mathbb{R}$, we let T_g^t be the hyperbolic isometry stabilizing g and acting by oriented translation distance t along g . The opposite orientation of g will be denoted \bar{g} , so that $T_g^t = T_{\bar{g}}^{-t}$.

For an oriented spike $\vec{s} = (g_u^v, g_u^w)$, its opposite orientation is $\bar{\vec{s}} = (\bar{g}_u^w, \bar{g}_u^v)$. In particular, we note that if \vec{s} is an oriented spike of H_u crossed by $k_{v,w}$, then $\bar{\vec{s}}$ is an oriented spike crossed by $k_{w,v} = \bar{k}_{v,w}$.

Shape-shifting in spikes. Suppose (v, w) is a simple pair and suppose u is between v and w . Let $\vec{s} = (g_u^v, g_u^w)$ be the spike of u crossed by $k_{v,w}$ with basepoints p_v and p_w . We define the elementary *shaping transformation* $A(\vec{s}) \in \operatorname{Isom}^+(\tilde{X}) = \operatorname{PSL}_2 \mathbb{R}$ determined by X , \mathfrak{s} , and s to be

$$(32) \quad A(\vec{s}) := T_{g_u^v}^{f_{X,s}(\vec{s})} \circ T_{g_u^w}^{-f_{X,s}(\vec{s})}.$$

Ultimately, the element $A(\vec{s})$ will be the value of the shear-shape cocycle $\varphi_{\mathfrak{s}}$ on the pair (g_u^v, g_u^w) ; see just below for an explanation of how we think of $A(\vec{s})$ as “changing the shape” of s .

Observe that $A(\vec{s})$ is a parabolic transformation preserving the common ideal endpoint of s . A familiar computation shows that in the spike determined by g_u^v and $A(\vec{s})g_u^w$, the orthogeodesics emanating from p_v and $A(\vec{s})p_w$ meet at a point distance $h_{\underline{A}+\mathfrak{a}}(s)$ from each (supposing that the deformation is small enough).

To the oriented spike \vec{s} of u , we also associate the *elementary shape-shift*

$$(33) \quad \varphi(\vec{s}) := T_{g_u^v}^{\mathfrak{s}(v,u)} \circ A(\vec{s}) \circ T_{g_u^w}^{-\mathfrak{s}(v,u)} = T_{g_u^v}^{\mathfrak{s}(v,u) + f_{X,s}(\vec{s})} \circ T_{g_u^w}^{-(\mathfrak{s}(v,u) + f_{X,s}(\vec{s}))}$$

where we recall that the value $\mathfrak{s}(v, u)$ is obtained by thinking of \mathfrak{s} as a function on transverse pairs (à la Lemma 13.11). Note that $\varphi(s)$ depends on our reference point v : whereas $A(\vec{s})$ is eventually identified as a value of the shape-shifting cocycle $\varphi_{\mathfrak{s}}$, the elementary shape-shifts $\varphi(\vec{s})$ are only building blocks for values of $\varphi_{\mathfrak{s}}$.

For the opposite orientation $\bar{s} = (\bar{g}_u^w, \bar{g}_u^v)$, we check

$$(34) \quad A(\bar{s}) = T_{\bar{g}_u^w}^{f_{X,\mathfrak{s}}(\bar{s})} T_{\bar{g}_u^v}^{-f_{X,\mathfrak{s}}(\bar{s})} = T_{g_u^w}^{f_{X,\mathfrak{s}}(\bar{s})} T_{g_u^v}^{-f_{X,\mathfrak{s}}(\bar{s})} = A(\vec{s})^{-1}.$$

Since $\mathfrak{s}(v, u) = \mathfrak{s}(u, v)$, we may similarly observe that $\varphi(\bar{s}) = \varphi(\vec{s})^{-1}$.

Take $\mathcal{H}_{v,w}$ to be the set of hexagons between v and w equipped with the linear order $u_1 < u_2$ induced by the orientation of $k_{v,w}$. Let $\underline{\mathcal{H}} \subset \mathcal{H}_{v,w}$ be any finite subset and order its elements $\underline{\mathcal{H}} = \{u_1, \dots, u_n\}$. For short, we denote hexagons $H_i := H_{u_i}$, spikes $s_i := \vec{s}_{u_i}$, geodesics $g_i^v := g_{u_i}^v$, etc.

To the finite, ordered set $\underline{\mathcal{H}}$ we associate the product

$$(35) \quad \varphi_{\underline{\mathcal{H}}} := \varphi(s_1) \circ \dots \circ \varphi(s_n) \circ T_{g_w^v}^{\mathfrak{s}(v,w)} \in \text{Isom}^+(\tilde{X}).$$

The goal of the rest of the section is then to extract a meaningful limit from $\varphi_{\underline{\mathcal{H}}}$ as $\underline{\mathcal{H}}$ increases to $\mathcal{H}_{v,w}$. Ultimately, this limit is how we will define the shape-shifting cocycle $\varphi_{\mathfrak{s}}$ on the boundary geodesics g_w^w and g_w^v corresponding to the simple pair (v, w) .

Remark 14.8. In the case that λ is maximal, each H_i is an ideal triangle and so $A = A + \mathfrak{a} = \emptyset$. In this case, each spike parameter $f_{X,\mathfrak{s}}(s_i)$ is 0 and we recover the formula from [Bon96], p. 255].

Geometric explanation of (33). Before proving convergence, however, let us explain the intuition behind the formulas above. In order to interpret $A(\vec{s})$ and $\varphi(\vec{s})$ as deformations of the hyperbolic structure X , we will switch our viewpoint to think of them as values of a deformation cocycle, and so as affecting the placement of pointed geodesics relative to each other. For brevity, let $f_{X,\mathfrak{s}}(\vec{s}) = t$.

Let us focus first on the shaping transformation $A(\vec{s})$. The oriented spike \vec{s} in the hexagon H_u is formed by two pointed geodesics (g_u^v, p_u^v) and (g_u^w, p_u^w) . Fixing our viewpoint at (g_u^v, p_u^v) , we may think of $A(\vec{s})$ as deforming \tilde{X} by holding (g_u^v, p_u^v) fixed and identifying (g_u^w, p_u^w) with $A(\vec{s}) \cdot (g_u^v, p_u^v)$. This has the overall effect of “widening” the spike s so that its sharpness parameter increases from h_A to $h_{A+\mathfrak{a}}$.

If instead we fix our basepoint to be outside of H_u , say at the basepoint p_v^w on $g_v^w \subset \partial_\lambda H_v$, then this transformation can be viewed as a composition of left and right earthquakes. Let Q^w and Q^v denote the half-spaces to the left of the oriented geodesics g_u^w and g_u^v , respectively. Note that $Q^w \subset Q^v$. The deformation $A(\vec{s})$ may then be thought of as first transforming all geodesics of \tilde{X} that lie in Q^w by $T_{g_u^w}^{-t}$; this has the effect of breaking \tilde{X} open along g_u^w and sliding Q_w to the left by distance t along g_u^w while keeping $\tilde{X} \setminus Q^w$ fixed. The deformation then further transforms all geodesics in Q^v by $T_{g_u^v}^t$; this is equivalent to the right earthquake with fault locus g_u^v that slides Q^v to the right while keeping $\tilde{X} \setminus Q^v$ fixed. The cumulative effect is then that the spike s has been “pushed” in the direction of \vec{s} by distance t . See Figure 22.

Remark 14.9. We give one final interpretation of $A(\vec{s})$ as “sliding G_u along H_u ” in the proof of Proposition 14.19 below (see also Figure 26), once we have set up the framework to understand the utility of this viewpoint.

In particular, note that the shear from H_v to H_u measured from p_v^w to the image of p_u^v under this composition of earthquakes has increased by $t = f_{X,\mathfrak{s}}(\vec{s})$. Therefore, if we let q_u^v denote the basepoint on g_u^v corresponding to the hexagon G_u , then the shear from H_v to H_u measured from p_v^w to the image of q_u^v under the deformation is exactly the original shear $\sigma_\lambda(X)(v, u)$ between v and u .

The elementary shape-shift $\varphi(\vec{s})$ can be interpreted in much the same way, but now the spike should be pushed distance $f_{X,\mathfrak{s}}(\vec{s}) + \mathfrak{s}(v, u)$ so that the resulting shear (measured between p_v^w and the image of q_u^v) is exactly $\sigma_\lambda(X)(v, u) + \mathfrak{s}(v, u)$.

Finally, the composition (35) can be thought of as a composition of the operations described above (read from right to left). Therefore, $\varphi_{\underline{\mathcal{H}}}$ first performs a right earthquake along g_w^v by $\mathfrak{s}(v, w)$, then performs an elementary shape-shift to pushing the spike s_n by $\mathfrak{s}(v, u_n) + f_{X,\mathfrak{s}}(s_n)$, then performs a shape-shift for s_{n-1} , etc. Observe that if q_i^v denotes the basepoint in g_i^v corresponding to G_{u_i} , then by construction the shear between v and each u_i measured from p_v^w to the image of q_i^v under the composite deformation is exactly the desired shear $\sigma_\lambda(X)(v, u_i) + \mathfrak{s}(v, u_i)$.

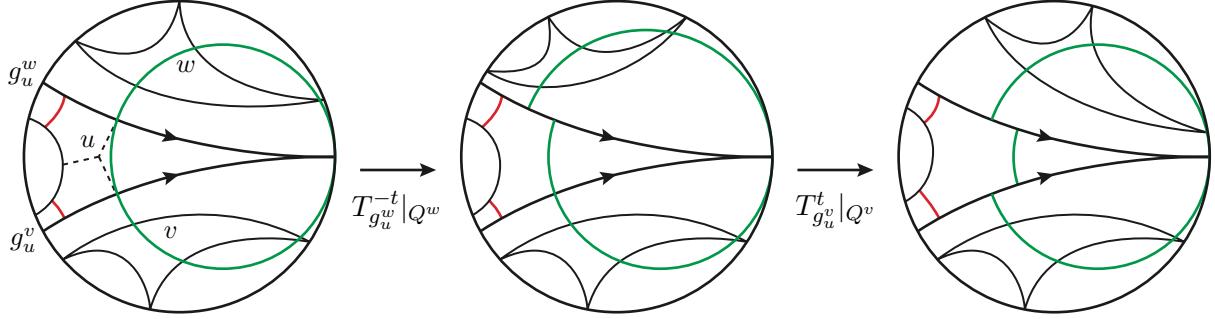


FIGURE 22. The effect of $A(\vec{s})$ when considered as a composition of left and right earthquakes.

Assuming the convergence of $\varphi_{\mathcal{H}}$ to a limit $\varphi_{v,w}$ (a step performed just below), we see that the placement of $\varphi_{\mathcal{H}}(g_w^v, p_w^v)$ limits to that of $\varphi_{v,w}(g_w^v, p_w^v)$. This in turn will be the placement of the geodesic (g_w^v, p_w^v) relative to (g_v^w, p_v^w) straightened in the deformed surface \tilde{X}_s ; see Lemma 15.6.

Convergence. We now consider the limiting behavior of $\varphi_{\mathcal{H}}$ as $\mathcal{H} \rightarrow \mathcal{H}_{v,w}$; that a limit exists is almost exactly the content of [Bon96, Lemma 14]. We give a proof here for convenience of the reader and to make sure that we are extracting the correct radius of convergence, i.e., that the modifications in the cusps actually do not affect the radius of convergence (even though there are countably many contributions from changing the shape of each spike).

Recall from Lemma 14.5 that the function $D_\lambda(X) = \text{inj}_\lambda(Y)/9|\chi(S)|$ gives a bound for the rate of decay of the length of a piece of a leaf of $\mathcal{O}_\lambda(X)$ in terms of its divergence radius.

Lemma 14.10 (compare Lemma 14 of [Bon96]). If $\|\mathfrak{s}\|_{\tau_\alpha} < D_\lambda(X)$, then $\varphi_{\mathcal{H}}$ converges to a well-defined isometry $\varphi_{v,w}$ as \mathcal{H} tends to $\mathcal{H}_{v,w}$.

Definition 14.11. The limiting isometry $\varphi_{v,w}$ is called the *shape-shifting map* for the simple pair (v, w) .

Remark 14.12. After combining all of our deformations in Section 14.5, the shape-shifting map $\varphi_{v,w}$ will be identified as the value of the shape-shifting cocycle $\varphi_{\mathfrak{s}}$ on the pair (g_w^v, g_w^v) . However, due to the asymmetry of our current definition, it is not clear that $\varphi_{v,w}^{-1} = \varphi_{w,v}$. See Lemma 14.14.

Proof of Lemma 14.10. For brevity, we set $D = D_\lambda(X)$ for the remainder of the proof.

Identify \tilde{X} with \mathbb{H}^2 and $\text{Isom}^+(\tilde{X})$ with the unit tangent bundle $T^1\mathbb{H}^2$ so that the identity I is the vector over $p_v \in \tilde{X}$ that is tangent to g_v^w and pointed in the positive direction with the orientation on g_v^w induced by $k_{v,w}$. Equip $T^1\mathbb{H}^2$ with a left-invariant metric d that is right-invariant with respect to the stabilizer of p_v . Finally, for $A \in \text{Isom}^+(\tilde{X})$, let $\|A\| := d(I, A)$, so that $\|AB\| \leq \|A\| + \|B\|$ holds by the triangle inequality.

We first show that $\varphi_{\mathcal{H}}$ stays in a compact set in $\text{Isom}^+(\tilde{X})$. Using boundedness, we then show that any sequence $\mathcal{H} \rightarrow \mathcal{H}$ is in fact Cauchy with respect to d , hence converges.

We start by bounding the lengths of segments of the form $k_{v,w} \cap H_u$, where $u \in \mathcal{H}_{v,w}$. To this end, construct a geometric train track τ from λ on X , and assume that the projection of $k_{v,w}$ meets τ transversely. Subdivide $k_{v,w}$ into arcs k_1, \dots, k_m whose projections meet τ once in branches b_1, \dots, b_m . For all but finitely many $u \in \mathcal{H}_{v,w}$, we have $k_{v,w} \cap H_u \subset k_j \setminus \tilde{\lambda}$ for some $j = 1, \dots, m$.

If $d \subset k_j \setminus \tilde{\lambda}$, we set $r(d)$ to be the depth $r_{b_j}(d)$ of d with respect to b_j and $r(d) = 1$, otherwise. By Lemma 14.5, there is $B > 0$ such that for all $u \in \mathcal{H}_{v,w}$,

$$\ell(k_{v,w} \cap H_u) \leq B e^{-Dr(k_{v,w} \cap H_u)}.$$

With this estimate in mind, our next task is to give a uniform bound on $\|\varphi_{\mathcal{H}} \circ T_{g_w^v}^{-s(v,w)}\|$ for all finite $\mathcal{H} \subset \mathcal{H}_{v,w}$. For each i , let $\gamma_i \in \text{Isom}^+(\tilde{X})$ be the isometry corresponding to the tangent vector over $k_{v,w} \cap g_i^v$ pointing toward the positive endpoint of g_i^v . Unpacking definitions, we may therefore write the shape-shift

$\varphi(s_i)$ as

$$\varphi(s_i) = \gamma_i T_{g_v^w}^{\mathfrak{s}(v, u_i) + f_{X, \mathfrak{s}}(s_i)} T_{h_i}^{-(\mathfrak{s}(v, u_i) + f_{X, \mathfrak{s}}(s_i))} \gamma_i^{-1},$$

where $h_i := \gamma_i^{-1} g_i^w$.

An explicit computation (in the upper half plane model, say) shows that

$$\begin{aligned} \left\| T_{g_v^w}^{\mathfrak{s}(v, u_i) + f_{X, \mathfrak{s}}(s_i)} T_{h_i}^{-(\mathfrak{s}(v, u_i) + f_{X, \mathfrak{s}}(s_i))} \right\| &\leq (e^{|\mathfrak{s}(v, u_i) + f_{X, \mathfrak{s}}(s_i)|} - 1) \ell(k_{v, w} \cap H_i) \\ &\leq B e^{|\mathfrak{s}(v, u_i) + f_{X, \mathfrak{s}}(s_i)| - Dr(k_{v, w} \cap H_i)}. \end{aligned}$$

By Lemma 14.3 and the triangle inequality, we have that

$$|\mathfrak{s}(v, u_i) + f_{X, \mathfrak{s}}(s_i)| \leq \|\mathfrak{s}\|_{\tau_\alpha} r(k_{v, w} \cap H_i) + \|\mathfrak{s}\|_{\bar{s}}$$

and so we conclude that

$$\|\gamma_i^{-1} \varphi(s_i) \gamma_i\| \leq B' e^{r(k_{v, w} \cap H_i)(\|\mathfrak{s}\|_{\tau_\alpha} - D)}$$

for $B' = B e^{\|\mathfrak{s}\|_{\bar{s}}}$.

Notice now that conjugation by γ_i changes the reference point of our calculation at a distance in the plane at most $\ell(k_{v, w})$, so the effect of $\gamma_i^{-1} \varphi(s_i) \gamma_i$ on $g_i^w \cap k_{v, w}$ is a displacement by $e^{\ell(k_{v, w})}$ times the quantity indicated above. Since this is independent of $\underline{\mathcal{H}}$, we have that

$$(36) \quad \|\varphi(s_i)\| = O\left(e^{r(k_{v, w} \cap H_i)(\|\mathfrak{s}\|_{\tau_\alpha} - D)}\right)$$

for any spike s_i corresponding to any hexagon u between v and w .

Expanding out $\varphi_{\underline{\mathcal{H}}}$ in terms of the $\varphi(s_i)$ (see (35)), we have that

$$\left\| \varphi_{\underline{\mathcal{H}}} \circ T_{g_w^v}^{-\mathfrak{s}(v, w)} \right\| = \left\| \prod_{i=1}^n \varphi(s_i) \right\| \leq \sum_{i=1}^n \|\varphi(s_i)\| = O\left(\sum_{i=1}^n e^{r(k_{v, w} \cap H_i)(\|\mathfrak{s}\|_{\tau_\alpha} - D)}\right).$$

Since there are a uniformly bounded number of gaps with given depth (Lemma 14.2), the last expression is bounded by the sum of at most $6|\chi(S)|$ many geometric series which are convergent so long as $\|\mathfrak{s}\|_{\tau_\alpha} < D$. Therefore, we see there is a compact set K in $\text{Isom}^+(\tilde{X})$ so that $\varphi_{\underline{\mathcal{H}}} \in K$ for any finite subset $\underline{\mathcal{H}} \subset \mathcal{H}_{v, w}$.

Now that we have shown the family of isometries $\{\varphi_{\underline{\mathcal{H}}}\}$ to be uniformly bounded, we can show that any sequence of refinements is in fact Cauchy. So suppose that $\underline{\mathcal{H}}_n$ increases to $\mathcal{H}_{v, w}$ and $|\underline{\mathcal{H}}_n| = n$. By construction, we may therefore write

$$\varphi_n = \psi \psi' \text{ and } \varphi_{n+1} = \psi \varphi(s_u) \psi',$$

where $\underline{\mathcal{H}}_{n+1} = \underline{\mathcal{H}}_n \cup \{u\}$ and $\psi, \psi' \in K$. Writing $\varphi_{n+1} = \psi \psi' \varphi(s_u) [\varphi(s_u)^{-1}, \psi'^{-1}]$, we have that

$$d(\varphi_n, \varphi_{n+1}) = \left\| \varphi(s_u) [\varphi(s_u)^{-1}, \psi'^{-1}] \right\|.$$

The zeroth order term in the Taylor expansion near the identity for the function $X \mapsto \|[X, \psi'^{-1}]\|$ is 0, because $[I, \psi'^{-1}] = I$. Since ψ'^{-1} stays in a compact set,

$$\left\| [\varphi(s_u)^{-1}, \psi'^{-1}] \right\| = O(\|\varphi(s_u)\|)$$

(see [Thu97, Theorem 4.1.6] or [Gel14, Lemma 1.1 of Lecture 2]).

Combining this estimate with the triangle inequality and (36), we get that

$$d(\varphi_n, \varphi_{n+1}) = O(\|\varphi(s_u)\|) = O\left(e^{r(k_{v, w} \cap H_u)(\|\mathfrak{s}\|_{\tau_\alpha} - D)}\right).$$

Now there are at most $6|\chi(S)|$ many $u \in \mathcal{H}_{v, w}$ with $r(k_{v, w} \cap H_u) = r$ (Lemma 14.2), so as $n \rightarrow \infty$ we must have that $r \rightarrow \infty$, and hence $d(\varphi_n, \varphi_{n+1}) \rightarrow 0$. Moreover, since this goes to 0 exponentially quickly, the sequence is in fact Cauchy. This completes the proof of the Lemma. \square

Shape-shifting as a limit of signed earthquakes. Here we give a different description of the shape-shifting map which forgoes approximations by “pushing spikes” in favor of approximations by left and right simple earthquakes (compare [EM06, Section III]). While this reformulation is symmetric and geometrically meaningful, it comes at the cost of restricting which approximating sequences $\underline{\mathcal{H}} \rightarrow \mathcal{H}$ actually yield convergent sequences of deformations $\varphi_{\underline{\mathcal{H}}}$. See also the remark at the top of page 261 in [Bon96].

Let (v, w) be a simple pair and fix a geometric train track τ snugly carrying λ . So long as τ is built from a small enough neighborhood, we may assume that the geodesic $k_{v,w}$ is transverse to the branches of τ . Then for each integer $r \geq 0$, let $\mathcal{H}_{v,w}^r$ denote the set of hexagons such that $k_{v,w} \cap H_u$ has depth at most r with respect to the branches of τ . Order

$$\mathcal{H}_{v,w}^r = (u_0 = v, u_1, \dots, u_n, u_{n+1} = w),$$

and for each $i = 0, \dots, n$, choose a geodesic h_i that separates the interior of H_i from the interior of H_{i+1} . Orient each h_i so that it crosses $k_{v,w}$ from left to right and set

$$(37) \quad \varphi_{v,w}^r = T_{h_0}^{\mathfrak{s}(u_0, u_1)} \circ A(s_1) \circ T_{h_1}^{\mathfrak{s}(u_1, u_2)} \circ A(s_2) \circ \dots \circ A(s_n) \circ T_{h_n}^{\mathfrak{s}(u_n, u_{n+1})}.$$

We now wish to show that $\varphi_{v,w}^r \rightarrow \varphi_{v,w}$ as $r \rightarrow \infty$. As in the case of $\varphi_{\underline{\mathcal{H}}} \rightarrow \varphi_{v,w}$, this argument will parallel that of [Bon96], with the extra complicating factor of the adjustments $A(s_i)$ to the shape of cusps.

The interpretation of (37) as a deformation cocycle is now similar to that of (35), but is now a combination of spike-shaping transformations together with simple earthquakes.

Let us give a description of the action of this deformation on the pointed geodesic (g_v^w, p_v^w) in $\partial_\lambda H_w$ closest to v . Reading the formula from right to left, we can obtain $\varphi_{v,w}^r(g_v^w, p_v^w)$ by first breaking \tilde{X} along $h_n = g_w^v$ and sliding the closed half-space containing H_w signed distance $\mathfrak{s}(u_n, u_{n+1})$, keeping the open half-space containing H_v fixed. Applying the spike shaping transformation $A(s_n)$ then preserves the natural basepoints p_n^v and p_n^w but increases the sharpness parameter $h_{\underline{A}}(s_n)$, making it match that of the spike in the hexagon G_{u_n} in the deformed metric $\underline{A} + \mathfrak{a}$. We then simply continue moving from w to v (i.e., backwards along $k_{v,w}$), alternating between signed earthquakes in the h_i and shaping the next spike until we reach g_v^w . Note that unlike the deformations associated to $\varphi_{\underline{\mathcal{H}}}$, each step of the process requires only local information about the spike s_i and the shear between u_i and u_{i+1} .

Lemma 14.13 (Lemma 16 of [Bon96]). So long as $\|\mathfrak{s}\|_{\tau_{\underline{\alpha}}} < D_\lambda(X)/2$, we have that $\lim_{r \rightarrow \infty} \varphi_{v,w}^r = \varphi_{v,w}$.

Proof. Using additivity, $\mathfrak{s}(u_i, u_{i+1}) = \mathfrak{s}(v, u_{i+1}) - \mathfrak{s}(v, u_i)$, we observe that

$$(38) \quad \varphi_{v,w}^r = \left(T_{h_0}^{\mathfrak{s}(v, u_1)} A(s_1) T_{h_1}^{-\mathfrak{s}(v, u_1)} \right) \left(T_{h_1}^{\mathfrak{s}(v, u_2)} A(s_2) T_{h_2}^{-\mathfrak{s}(v, u_2)} \right) \dots \left(T_{h_{n-1}}^{\mathfrak{s}(v, u_n)} A(s_n) T_{h_n}^{-\mathfrak{s}(v, u_n)} \right) T_{h_n}^{\mathfrak{s}(v, w)}.$$

So $\varphi_{v,w}^r$ is obtained from $\varphi_{\mathcal{H}_{v,w}^r}$ by replacing each term of the form

$$\varphi(s_i) = T_{g_i^v}^{\mathfrak{s}(v, u_i)} A(s_i) T_{g_i^w}^{-\mathfrak{s}(v, u_i)}$$

with

$$\phi(s_i) := T_{h_{i-1}}^{\mathfrak{s}(v, u_i)} A(s_i) T_{h_i}^{-\mathfrak{s}(v, u_i)}$$

and $T_{g_w^v}^{\mathfrak{s}(v, w)}$ with $T_{h_n}^{\mathfrak{s}(v, w)}$.

The basic estimate we need is approximately how close $\varphi(s_i)$ is to $\phi(s_i)$ in $\text{Isom}^+(\tilde{X})$ as r tends to infinity. For this we will want to understand how closely h_{i-1} approximates g_i^v near its intersection with $k_{v,w}$, as well as for g_i^w and h_i .

By construction, h_i must be between g_i^w and g_{i+1}^v for each $i = 1, \dots, n$ and h_0 is between g_v^w and g_1^v . But g_i^w and g_{i+1}^v follow the same edge path of length $2r$ in $\tau \subset \tau_{\underline{\alpha}}$, for otherwise, there would be a another $u \in \mathcal{H}_{v,w}^r$ such that H_u separates H_i from H_{i+1} . Thus h_i follows the same edge path and fellow travels g_i^w and g_{i+1}^v for length at least $O(2rD_\lambda(X))$; using negative curvature, we have that h_i is $O(e^{-D_\lambda(X)r})$ close to both g_i^w and g_{i+1}^v near $k_{v,w}$.

From closeness of these geodesics from the previous paragraph (and our estimates for $\|\varphi(s_i)\|$ from Lemma 14.10) it is possible to obtain the basic estimate

$$\|\phi(s_i)^{-1} \varphi(s_i)\| = O(\exp(\|\mathfrak{s}\|_{\tau_{\underline{\alpha}}} r(k_{v,w} \cap H_i) - rD_\lambda(X))),$$

which is small when $\|\mathfrak{s}\|_{\tau_{\underline{\alpha}}} < D_\lambda(X)$. Notice that we have also used the fact that the adjustment parameter associated to each spike $\mathfrak{s}(s_i)$ is uniformly bounded; that said, even if it grew linearly in r we would obtain the same estimate (up to a multiplicative factor).

The rest of the argument ensuring that $\varphi_{v,w}^r$ and $\varphi_{\mathcal{H}_{v,w}^r}$ have the same limit as long as $\|\mathfrak{s}\|_{\tau_{\underline{\alpha}}} < D_\lambda(X)/2$ follows [Bon96, Lemma 16] and is omitted. We remark that the factor of $1/2$ appearing at the end is a relic of the techniques used in [Bon96, Lemma 16]. \square

The following simple fact was not apparent from the definition of $\varphi_{v,w}$ due to its lack of symmetry. Fortunately, the approximation of $\varphi_{v,w}$ by $\varphi_{v,w}^r$ gives us a symmetric description of $\varphi_{v,w}$.

Corollary 14.14. If (v,w) is simple and $\|\mathfrak{s}\|_{\tau_{\underline{\alpha}}} < D_\lambda(X)/2$, then

$$\varphi_{w,v} = \varphi_{v,w}^{-1}.$$

Proof. We observe first that $\mathcal{H}_{v,w}^r = \mathcal{H}_{w,v}^r$, so the each term of $\varphi_{v,w}^r$ appears in $\varphi_{w,v}^r$ with the opposite orientation. Now by [34], the inverse of the shaping transformation of an oriented spike is equal to the shaping transformation of the same spike with opposite orientation. Therefore we have that $\varphi_{v,w}^r = (\varphi_{w,v}^r)^{-1}$ for all r , and the equality holds as we take $r \rightarrow \infty$. \square

14.3. Shape-shifting in hexagons. In this section, we explain how to define the shape-shifting cocycle $\varphi_{\mathfrak{s}}$ on pairs of basepointed geodesics that lie in the boundary of a common hexagon; this will encode the change in hyperbolic structure on $X \setminus \lambda$.

While in this setting we do not have to worry about delicate convergence results, we must be more diligent about recording the placement of basepoints on each geodesic of $\partial_\lambda H_u$. Moreover, the cocycle condition (Propositions 14.18 and 14.19) only becomes apparent once we reinterpret the shaping deformations defined below as “sliding the deformed hexagon along the original.”

Throughout this section, we have extended both \underline{A} and $\underline{A} + \mathfrak{a}$ to some common maximal arc system $\underline{\alpha}$ by adding in arcs of weight 0 as necessary. We remind the reader that $\mathfrak{s}(\alpha)$ denotes the coefficient of α in \mathfrak{a} .

Notations and orientations. Let $H_u \subset \tilde{X} \setminus \tilde{\lambda}_{\underline{\alpha}}$ be a nondegenerate right-angled hexagon and enumerate the λ -boundary components of H_u as $\partial_\lambda H_u = \{(h_1, p_1), (h_2, p_2), (h_3, p_3)\}$, cyclically ordered about u . Let $\alpha_i \in \underline{\alpha}$ be the orthogeodesic arc opposite to h_i , and denote by p_{ij} the vertex of H_u meeting both h_i and α_j . See Figure 23. If H_u is a degenerate hexagon (i.e., a pentagon with one ideal vertex or a quadrilateral with two) then we label only those points and geodesics which appear in its boundary.

Each choice of orientation $\vec{\alpha}_1$ of α_1 induces orientations of h_2 and h_3 so that α_1 leaves from the left-hand side of h_j and arrives on the right-hand side of h_k for $\{j, k\} = \{2, 3\}$; an example is pictured in Figure 23. Observe that the opposite orientation $\vec{\alpha}_1$ induces the opposite orientations on h_2 and h_3 . Throughout this section, we also adopt similar conventions for each orientation of α_2 and α_3 .

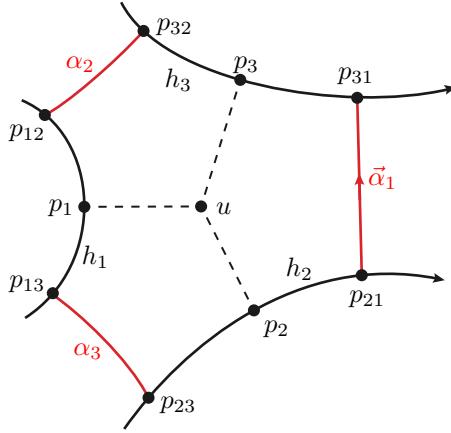


FIGURE 23. Distinguished points on a hexagon H_u and induced orientations on $h_2, h_3 \in \partial_\lambda H_u$.

Recall that (by Theorem 6.4) the deformation \mathfrak{s} induces a new metric on $\tilde{X} \setminus \tilde{\lambda}$ denoted by $\underline{A} + \mathfrak{a}$ and which contains a hexagon G_u corresponding to H_u . The corresponding basepointed λ -boundary geodesics and vertices of G_u will be denoted by (g_i, q_i) and q_{ij} , respectively. We adopt similar orientation conventions as above for the realizations of α_j in G_u .

Shapes of Hexagons. Paralleling our discussion for spikes, we first need to define geometric parameters that measure the shape of the hexagon as well as the difference of the placements of the basepoints p_i and

q_i on the geodesics h_i and g_i . For concreteness, we only consider α_1 below; the parameters for α_2 and α_3 are defined symmetrically.

We begin by associating to α_1 the parameter

$$\ell_s(\alpha_1) := \ell_{\underline{A}+\mathfrak{a}}(\alpha_1) - \ell_{\underline{A}}(\alpha_1) \in \mathbb{R}.$$

which measures the difference in the hyperbolic length of α_1 in the metric determined by $\underline{A} + \mathfrak{a}$ versus in the original metric \underline{A} induced by X .

Now fix an orientation $\vec{\alpha}_1$ of α_1 ; as above, this induces orientations of the geodesics h_2 , h_3 , g_2 , and g_3 . Let $d_{\underline{A}}(\vec{\alpha}_1, u)$ be the signed distance from p_2 to p_{21} on h_2 ; ²¹ the local symmetry of the orthogeodesic foliation implies that $d_{\underline{A}}(\vec{\alpha}_1, u)$ can also be computed as the signed distance from p_3 to p_{31} on h_3 . Define similarly $d_{\underline{A}+\mathfrak{a}}(\vec{\alpha}_1, u)$ as the distance from q_2 to q_{21} on g_2 (equivalently, the signed distance from q_3 to q_{31} on g_3).

To all of this information, we associate the parameter

$$f_{X,s}(\vec{\alpha}_1, u) := d_{\underline{A}+\mathfrak{a}}(\vec{\alpha}_1, u) - d_{\underline{A}}(\vec{\alpha}_1, u) \in \mathbb{R}$$

which measures the difference in how far u is from α_1 in G_u versus in H_u . More precisely, considering H_u and G_u in the hyperbolic plane, we can use an element of $\mathrm{PSL}_2(\mathbb{R})$ to line up (h_2, p_{21}) with (g_2, q_{21}) so that the basepoints and orientations agree. The parameter $f_{X,s}(\vec{\alpha}_1, u)$ then measures the distance from q_2 to p_2 along $h_2 = g_2$. See Figure 24. Of course, symmetry shows that it is equivalent to align (h_3, p_{31}) with (g_3, q_{31}) and measure the signed distance from q_3 to p_3 along $h_3 = g_3$.

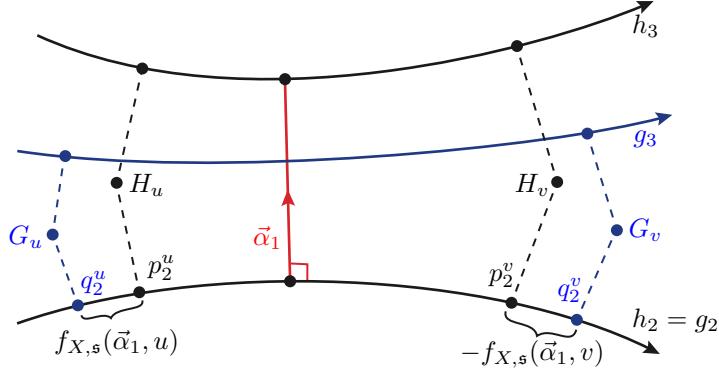


FIGURE 24. The parameter $f_{X,s}(\vec{\alpha}_1, u)$ for two adjacent hexagons. We have decorated the basepoints on $h_2 = g_2$ with a superscript to emphasize their dependence on the hexagon.

Note that reversing orientations reverses signs, so that $d_{\underline{A}}(\vec{\alpha}_1, u) = -d_{\underline{A}}(\vec{\alpha}_1, u)$ and hence

$$f_{X,s}(\vec{\alpha}_1, u) = -f_{X,s}(\vec{\alpha}_1, v).$$

The parameters associated to the hexagons which border a given arc are related in the following way:

Lemma 14.15. Let α be any edge of $\tilde{\alpha}$ and let H_u and H_v be its adjoining hexagons. Then

$$f_{X,s}(\vec{\alpha}, u) + f_{X,s}(\vec{\alpha}, v) = s(\alpha)$$

where the orientation $\vec{\alpha}$ is chosen so that u is on its left (equivalently, $\vec{\alpha}$ is oriented so that v is on its left).

Proof. The proof is an exercise in unpacking the definitions and being careful with orientations; compare Figure 24. Let p_2^u and p_2^v denote the projections of u and v to either of geodesics common to $\partial_\lambda H_u$ and $\partial_\lambda H_v$, and let q_2^u and q_2^v play similar roles for G_u and G_v .

We can then write

$$\begin{aligned} f_{X,s}(\vec{\alpha}, u) + f_{X,s}(\vec{\alpha}, v) &= f_{X,s}(\vec{\alpha}, u) - f_{X,s}(\vec{\alpha}, v) \\ &= d_{\underline{A}+\mathfrak{a}}(\vec{\alpha}, u) - d_{\underline{A}}(\vec{\alpha}, u) - d_{\underline{A}+\mathfrak{a}}(\vec{\alpha}, v) + d_{\underline{A}}(\vec{\alpha}, v) \\ &= d_h(p_2^u, p_2^v) - d_g(q_2^u, q_2^v) = s(\alpha) \end{aligned}$$

²¹The parameter $d_{\underline{A}}(\vec{\alpha}_1, u)$ is called the “ t -coordinate” of the arc α_1 in the hexagon H_u in [Luo07]. See also [Mon09b, Proposition 2.10], where a formula is given in terms of the lengths of $\{\alpha_i : i = 1, 2, 3\}$.

where we recall that $\mathfrak{s}(\alpha)$ denotes the coefficient of α in \mathfrak{a} and where d_h and d_g represent the signed distance measured along h_2 and g_2 , equipped with the orientation induced by $\vec{\alpha}$. \square

Remark 14.16. Using Theorem 6.4 and some hyperbolic trigonometry, one may show that $f_{X,\mathfrak{s}}(\vec{\alpha}_1, u)$ depends analytically on both A and \mathfrak{a} for fixed α_1 and u .

Shaping Hexagons. To the hexagon H_u and oriented arc $\vec{\alpha}_1$ in its boundary, we associate the *shaping transformation* $A(\vec{\alpha}_1, u)$ by

$$(39) \quad A(\vec{\alpha}_1, u) := T_{h_2}^{-f_{X,\mathfrak{s}}(\vec{\alpha}_1, u)} \circ T_{\vec{\alpha}_1}^{\ell_{\mathfrak{s}}(\alpha_1)} \circ T_{h_3}^{f_{X,\mathfrak{s}}(\vec{\alpha}_1, u)} \in \text{Isom}^+(\widetilde{X}),$$

where $T_{\vec{\alpha}_1}$ denotes translation along the complete oriented geodesic extending $\vec{\alpha}_1$. The shaping transformation is explicitly constructed so that if H_u and G_u are superimposed with $(h_2, p_2) = (g_2, q_2)$, then

$$A(\vec{\alpha}_1, u)(h_3, p_3) = (g_3, q_3).$$

This claim is not immediately apparent from the expression of (39), but is easy to verify once we reinterpret $A(\vec{\alpha}_1, u)$ as “sliding G_u along H_u .”

To wit, suppose that we superimpose H_u and G_u so that $(h_3, p_3) = (g_3, q_3)$. Now consider what happens as we apply $A(\vec{\alpha}_1, u)$ to G_u while holding H_u fixed; the first term $T_{h_3}^{f_{X,\mathfrak{s}}(\vec{\alpha}_1, u)}$ translates G_u along h_3 so that $q_{31} = p_{31}$, and the right angle formed by α_1 and g_3 in G_u lines up with the same angle in H_u . The transformation $T_{\vec{\alpha}_1}^{\ell_{\mathfrak{s}}(\alpha_1)}$ then slides $T_{h_3}^{f_{X,\mathfrak{s}}(\vec{\alpha}_1, u)}G_u$ along α_1 so that $(h_2, q_{21}) = (g_2, p_{21})$. Finally, $T_{h_2}^{-f_{X,\mathfrak{s}}(\vec{\alpha}_1, u)}$ slides $T_{\vec{\alpha}_1}^{\ell_{\mathfrak{s}}(\alpha_1)}T_{h_3}^{f_{X,\mathfrak{s}}(\vec{\alpha}_1, u)}G_u$ along $h_2 = g_2$ so that q_2 lines up with p_2 . See Figure 25.

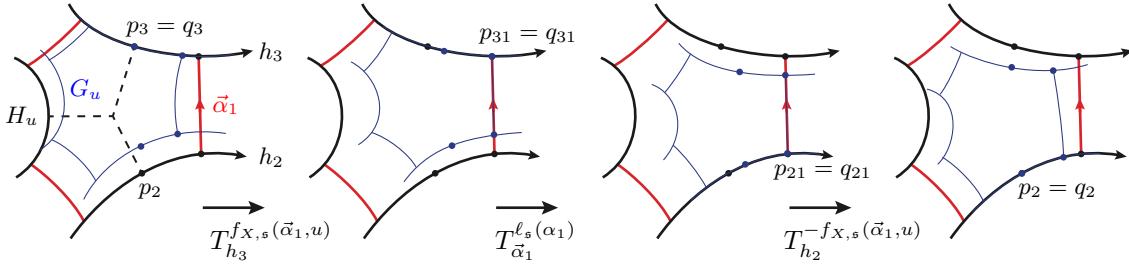


FIGURE 25. How the shaping transformation $A(\vec{\alpha}_1, u)$ slides G_u along H_u .

Summarizing, we have shown that $A(\vec{\alpha}_1, u)$ takes a superimposition of G_u on H_u with $(h_3, p_3) = (g_3, q_3)$ to another superimposition with $(h_2, p_2) = (g_2, q_2)$. In particular, this implies that applying $A(\vec{\alpha}_1, u)$ to (h_3, p_3) takes it to the position of (g_3, q_3) in the latter placement of G_u , which is what we claimed.

Remark 14.17. An elementary hyperbolic geometry argument similar to that in the proof of Lemma 14.6 shows that if α_1 in X degenerates to an oriented spike \vec{s} then the corresponding geometric parameter $f_{X,\mathfrak{s}}(\vec{\alpha}_1, u)$ limits to the parameter $f_{X,\mathfrak{s}}(\vec{s})$. In particular, along this degeneration the corresponding hexagon-shaping transformation $A(\vec{\alpha}_1, u)$ converges to the spike-shaping transformation $A(\vec{s})$.

A cocycle condition for hexagons. A number of relations hold between the shaping transformations for different arcs and different orientations; eventually, these relations are what ensure that the deformations we are currently building package together into an honest cocycle (see Proposition 14.26).

First, we observe that reversing the orientation of α_1 inverts the shaping transformation:

$$(40) \quad A(\vec{\alpha}_1, u) = T_{h_3}^{-f_{X,\mathfrak{s}}(\vec{\alpha}_1, u)} \circ T_{\vec{\alpha}_1}^{\ell_{\mathfrak{s}}(\alpha_1)} \circ T_{h_2}^{f_{X,\mathfrak{s}}(\vec{\alpha}_1, u)} = T_{h_3}^{-f_{X,\mathfrak{s}}(\vec{\alpha}_1, u)} \circ T_{\vec{\alpha}_1}^{-\ell_{\mathfrak{s}}(\alpha_1)} \circ T_{h_2}^{f_{X,\mathfrak{s}}(\vec{\alpha}_1, u)} = A(\vec{\alpha}_1, u)^{-1}.$$

Now suppose that H_u is on the left and H_v is on the right of the oriented arc $\vec{\alpha}_1$. Combining the relation of Lemma 14.15 with the definition of the shaping transformation, we have that

$$(41) \quad A(\vec{\alpha}_1, u) = T_{h_2}^{\mathfrak{s}(\alpha_1)} \circ A(\vec{\alpha}_1, v) \circ T_{h_3}^{-\mathfrak{s}(\alpha_1)}.$$

This equation is used frequently in Section 14.4 just below.

Finally, a beautiful and important relationship holds among the three shaping transformations in a single right-angled hexagon. Our proof utilizes the “sliding” viewpoint explained above; the statement seems difficult to prove just by writing down a string of Möbius transformations.

Proposition 14.18. Let $u \in \mathcal{H}$ be a nondegenerate right-angled hexagon with boundary arcs $\vec{\alpha}_1, \vec{\alpha}_2, \vec{\alpha}_3$, oriented so that H_u lies to the left of each $\vec{\alpha}_i$. Then

$$A(\vec{\alpha}_3, u) \circ A(\vec{\alpha}_2, u) \circ A(\vec{\alpha}_1, u) = 1.$$

A similar statement clearly holds for any cyclic permutation of $(3, 2, 1)$.

Proof. In order to prove the lemma, we superimpose G_u on top of H_u so that $(g_3, q_3) = (h_3, p_3)$. Holding H_u fixed, the first shaping transformation $A(\vec{\alpha}_1, u)$ slides G_u along h_3 , then along α_1 , then along h_2 so that (g_2, q_2) lines up with (h_2, p_2) . The second shaping transformation $A(\vec{\alpha}_2, u)$ then acts on this translated copy of G_u by sliding it along h_2 , then α_2 , then h_1 so that $(g_1, q_1) = (h_1, p_1)$. Finally, the last term slides $A(\vec{\alpha}_1, u)A(\vec{\alpha}_2, u)G_u$ along the edges of H_u so that (g_3, q_3) returns to (h_3, p_3) (with the same orientation).

Therefore, since $A(\vec{\alpha}_1, u) \circ A(\vec{\alpha}_2, u) \circ A(\vec{\alpha}_3, u)$ preserves a unit tangent vector and $\text{Isom}^+(\tilde{X})$ acts simply transitively on $T^1\tilde{X}$, the composition of the three shaping transformations must be trivial. \square

A similar result holds for degenerate right-angled hexagons, with the hexagon-shaping transformation replaced with the corresponding spike-shaping transformation.

Proposition 14.19. Suppose that $u \in \mathcal{H}$ is a pentagon with two orthogeodesic arcs α_1, α_2 and one spike s , labeled so that (α_1, α_2, s) runs counterclockwise around u . Orient each α_j so that H_u is on its left and orient s so that it is pointing towards the ideal vertex of H_u . Then

$$A(\vec{s}) \circ A(\vec{\alpha}_2, u) \circ A(\vec{\alpha}_1, u) = 1.$$

Similarly, if $u \in \mathcal{H}$ is a quadrilateral with one orthogeodesic edge α and two spikes s_1 and s_2 (labeled so that (α, s_1, s_2) is read counterclockwise), then

$$A(\vec{s}_2) \circ A(\vec{s}_1) \circ A(\vec{\alpha}, u) = 1$$

where all orientation conventions are as above.

Proof. We only explain how to interpret the spike-shaping transformation $A(\vec{s})$ in our “sliding” framework; once we have done so, the rest of the proof is completely analogous to that of Proposition 14.18.

So let \vec{s} be a spike of H_u , oriented as described; suppose that its left and right boundary geodesics are h_3 and h_2 . Recall that $A(\vec{s})$ is constructed so that if we superimpose G_u and H_u with $(g_2, q_2) = (h_2, p_2)$ then $A(\vec{s})(h_3, p_3) = (g_3, q_3)$. This can equivalently be interpreted by superimposing G_u on H_u with $(g_3, q_3) = (h_3, p_3)$; then applying the shaping transformation to G_u while leaving H_u fixed takes G_u to another superimposition where $(g_2, q_2) = (h_2, p_2)$. See Figure 26. \square

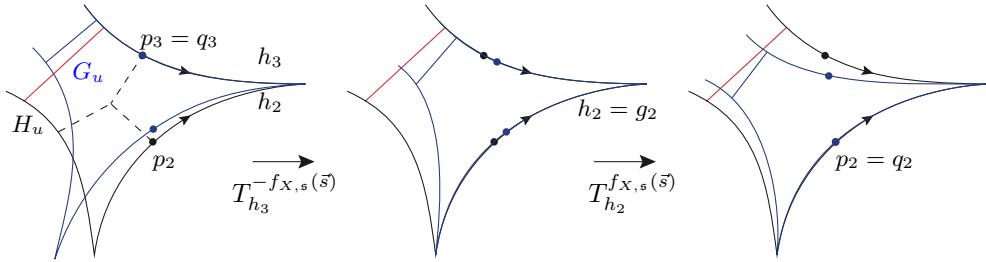


FIGURE 26. Interpreting the spike-shaping transformation $A(\vec{s})$ as sliding G_u along H_u . Note that in this picture, we have $f_{X,s}(\vec{s}) < 0$.

14.4. Shape-shifting along the spine. In this section we package together the hexagon-shaping deformations defined in (39) into deformations of entire complementary subsurfaces of $\tilde{X} \setminus \tilde{\lambda}$. As always, we will exhibit this deformation by explaining how to adjust the positions of the pointed geodesics in the boundary of each component of $\tilde{X} \setminus \tilde{\lambda}$ relative to one another. This in turn requires some book-keeping of orientations and a liberal application of the cocycle relation (Propositions 14.18 and 14.19).

Throughout this section, we fix some component Y of $\tilde{X} \setminus \tilde{\lambda}$. We remind the reader that the deformation \mathfrak{s} induces a new hyperbolic structure $\underline{A} + \mathfrak{a}$ on Y whose hexagons and basepointed geodesics correspond to those of Y .

Hexagonal hulls and induced orientations. Suppose that $v, w \in \mathcal{H}$ are distinct hexagons of Y . Since the corresponding component of $\tilde{\text{Sp}}$ is a tree it contains a unique oriented non-backtracking edge path $[v, w]$ joining v to w . We then define the *hexagonal hull* $H(v, w)$ of (v, w) to be the union of all of the hexagons corresponding to the vertices of $[v, w]$. Define also the *truncated hexagonal hull* $\widehat{H}(v, w)$ by truncating each spike of $H(v, w)$ by the horocycle through the basepoints that are closest to the ideal vertex. Note that both $H(v, w)$ and $\widehat{H}(v, w)$ come with $(\pi_1(Y)$ -equivariant) collections of basepoints on their boundaries obtained by projecting each of the vertices of $[v, w]$ onto the associated boundary geodesics.

Now for any $(h_v, p_v) \in \partial_\lambda H_v$ and $(h_w, p_w) \in \partial_\lambda H_w$, we have that $\partial \widehat{H}(v, w) \setminus \{p_v, p_w\}$ consists of two paths δ_\pm . We orient each of δ_\pm so that they both travel from p_v to p_w . See Figure 27.

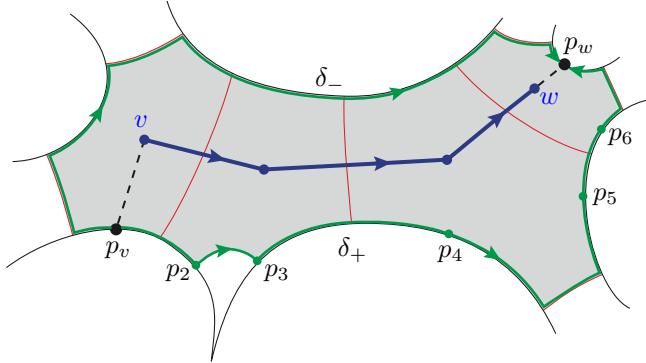


FIGURE 27. The truncated hexagonal hull (shaded) of the path $[v, w]$ and the induced orientations on the paths δ_\pm from p_v to p_w in its boundary.

With this induced orientation, the path δ_+ passes through a sequence of basepoints

$$p_v = p_1, p_2, \dots, p_{n+1} = p_w.$$

We then associate a shaping transformation A_i to each subsequent pair of basepoints as follows:

- If p_i, p_{i+1} are in different hexagons, then they must lie on the same geodesic h_i of ∂Y and correspond to two hexagons H_i and H_{i+1} both adjacent to an arc α_i . In this case, define $A_i = T_{h_i}^{\mathfrak{s}(\alpha_i)}$ where h_i is given the orientation induced by δ_+ and where we recall that $\mathfrak{s}(\alpha_i)$ is the coefficient of α_i in \mathfrak{a} .
- If p_i, p_{i+1} are in the same hexagon H_{u_i} but do not lie on a common spike, then necessarily they lie on geodesics connected by some arc α_i . In this case, define $A_i = A(\vec{\alpha}_i, u_i)$ where the orientation on α_i is induced from δ_+ .
- If p_i and p_{i+1} lie on a common spike s_i , then we define $A_i = A(\vec{s}_i)$, where the orientation on s_i is such that the horocyclic segment of δ_+ cutting off s_i runs from the left of one of the oriented geodesics to the right of the other.

Finally, we then combine all of this information to define the shape-shifting transformation

$$(42) \quad A(\delta_+) := A_1 \circ A_2 \circ \dots \circ A_n,$$

where we recall that we are multiplying from right to left. Define $A(\delta_-)$ analogously; the point is, however, that the choice of \pm does not matter.

Lemma 14.20. We have $A(\delta_-) = A(\delta_+)$.

Definition 14.21. We call $\varphi_{p_v, p_w} := A(\delta_+) = A(\delta_-)$ the *shape-shifting map* for the pair $((h_v, p_v), (h_w, p_w))$.

Proof. The proof follows by induction on the length of $[v, w]$. If $[v, w]$ has length 0, i.e., $v = w$, then this statement is exactly the content of the cocycle relation for hexagons (Propositions 14.18 and 14.19).

Now suppose that $[v, w]$ has length n and let u be the penultimate vertex in $[v, w]$. Let α denote the arc separating u from w , and choose the orientation $\vec{\alpha}$ so that u lies on its left. Up to relabeling, we may assume that the orientation of δ_+ agrees with the orientation of $\partial\widehat{H}(v, w)$. Denote by $(h_u^\pm, p_u^\pm) \in \partial Y$ the last basepoints of H_u visited by δ_\pm and let γ_\pm denote the subpaths of δ_\pm from p_v to p_u^\pm in the boundary of the truncated hexagonal hull $\widehat{H}(v, u)$. Define $A(\gamma_\pm)$ analogously to $A(\delta_\pm)$. Then we may write

$$\begin{aligned} A(\delta_+)A(\delta_-)^{-1} &= A(\gamma_+)T_{h_u^+}^{s(\alpha)}B_1B_2T_{h_u^-}^{-s(\alpha)}A(\gamma_-)^{-1} \\ &= (A(\gamma_+)A(\vec{\alpha}, u)A(\gamma_-)^{-1}) \cdot A(\gamma_-)(A(\vec{\alpha}, u)^{-1}T_{h_u^+}^{s(\alpha)}B_1B_2T_{h_u^-}^{-s(\alpha)})A(\gamma_-)^{-1} \end{aligned}$$

where B_1 and B_2 are the shaping transformations corresponding to arcs and spikes of w that are different from α (labeled counterclockwise from α), oriented either so that w lies on the left of the arc or so that the spike points into the common ideal endpoint.

Now observe that $A(\gamma_+)A(\vec{\alpha}, u)A(\gamma_-)^{-1}$ is trivial by the inductive hypothesis, as it corresponds to the comparison between the two possible definitions of φ_{p_v, p_u^-} . We also note that

$$A(\vec{\alpha}, u)^{-1}T_{h_u^+}^{s(\alpha)}B_1B_2T_{h_u^-}^{-s(\alpha)}$$

is conjugate to

$$T_{h_u^-}^{-s(\alpha)}A(\vec{\alpha}, u)^{-1}T_{h_u^+}^{s(\alpha)}B_1B_2 = A(\vec{\alpha}, w)B_1B_2 = 1$$

where the first equality follows from (41) (note the reversals in orientations of h_u^\pm) and the second follows from the cocycle relation (Propositions 14.18 and 14.19). Therefore, we see that the entire term $A(\delta_+)A(\delta_-)^{-1}$ is trivial, which is what we wanted to show. \square

Remark 14.22. The above statement can also be proven by interpreting $A(\delta_\pm)$ in terms of sliding. In particular, let Z denote the $\pi_1(Y)$ -equivariant hyperbolic structure on Y corresponding to the weighted arc system $\underline{A} + \underline{a}$. Then superimposing Z on Y so that $(g_w, q_w) = (h_w, p_w)$, one can consecutively apply the shaping transformations A_i to Z while keeping Y fixed.

Doing so, we see that A_n moves Z so that $(g_n, q_n) = (h_n, p_n)$, then $A_{n-1} \circ A_n$ moves Z so that $(g_{n-1}, q_{n-1}) = (h_{n-1}, p_{n-1})$, etc. At the end of this process, we have applied $A(\delta_+)$ to Z and by construction, the pointed geodesic (g_v, q_v) matches up with (h_v, p_v) . Since the final positioning of Z is the same relative to Y whether we used $A(\delta_+)$ or $A(\delta_-)$, this allows us to conclude that the two compositions define the same element.

Remark 14.23. While we used the distinguished boundary paths δ_\pm to define the shape-shifting map, one could in fact use *any* path from p_v to p_w in $Y \cup \widetilde{\alpha}$. In this case, one must take more care to enumerate basepoints so that p_i, p_{i+1} always either lie on the same geodesic or in the same hexagon.

Observe that reversing the orientation $\overline{[v, w]} = [w, v]$ also reverses the sequence p_{n+1}, \dots, p_1 of basepoints that the boundary paths $\overline{\delta_\pm}$ meet. Since flipping the order of p_i and p_{i+1} inverts each of the A_i transformations defined above, we therefore discover that $\varphi_{p_w, p_v} = \varphi_{p_v, p_w}^{-1}$.

In a similar vein, it is not hard to see that the shape-shifting maps satisfy a cocycle relation.

Proposition 14.24. For any triple of pointed geodesics (h_u, p_u) , (h_v, p_v) , and (h_w, p_w) of ∂Y , we have that

$$\varphi_{p_u, p_v} \circ \varphi_{p_v, p_w} \circ \varphi_{p_w, p_u} = 1.$$

Proof. This follows immediately from the definitions when v lies on either of the paths δ_\pm from u to w .

Otherwise, note that the intersection of paths $[u, v] \cap [v, w] \cap [w, u]$ is a point $x \in \mathcal{H}$. Choosing a basepoint $p_x \in \partial_\lambda H_x$, compute the shape-shifting transformations using the boundary arcs that pass through x . Then we may express $\varphi_{p_u, p_v} = \varphi_{p_u, p_x} \circ \varphi_{p_x, p_v}$ and using the observation about inverses above, we realize that

$$\varphi_{p_u, p_v} \circ \varphi_{p_v, p_w} \circ \varphi_{p_w, p_u} = \varphi_{p_u, p_x} \circ (\varphi_{p_x, p_v} \circ \varphi_{p_v, p_x}) \circ (\varphi_{p_x, p_w} \circ \varphi_{p_w, p_x}) \circ \varphi_{p_x, p_u} = 1.$$

This completes the proof of the proposition. \square

14.5. The shape-shifting cocycle. We now combine the shape-shifting maps for simple pairs (Definition 14.11) with those for complementary subsurfaces (Definition 14.21) into the promised shape-shifting cocycle (Proposition 14.26), which is well-defined as long as the combinatorial deformation \mathfrak{s} is small enough. As usual, we construct a geometric train track τ from λ on X so that the weight space of $\tau_{\underline{\alpha}}$ provides a notion of size for \mathfrak{s} .

Admissible routes. For $v, w \in \mathcal{H}$ and Y a component of $\tilde{X} \setminus \tilde{\lambda}$, we say that Y is *thick* with respect to v and w if one of the two possibilities occur:

- (1) either Y contains v and/or w , or
- (2) v and w lie in different components of $\tilde{X} \setminus Y$ and the boundary leaves of Y closest to v and w are not asymptotic.

Observe that in the first case, there is either no or one boundary geodesic of Y separating v from w (depending if v and w are both in Y or not), while in the second, there are exactly two boundary components of Y separating v from w .

Now let $v, w \in \mathcal{H}$ be any pair of distinct hexagons that do not lie in the same component of $\tilde{X} \setminus \tilde{\lambda}$ and let (h_v, p_v) and (h_w, p_w) be a pointed geodesic in $\partial_{\lambda} H_v$ and $\partial_{\lambda} H_w$. Then there is a unique (possibly empty) sequence h_1, \dots, h_n of boundary geodesics of thick subsurfaces separating p_v from p_w , ordered by proximity to v (with h_1 closest). ²² If one of the h_i lies in the boundary of two complementary subsurfaces (so corresponds to a lift of a curve component of λ) then we record it twice, one time for each of the adjoining subsurfaces. Additionally, if either h_v or h_w is a boundary geodesic separating v from w , then we do not record it as one of the h_i . See Figure 28.

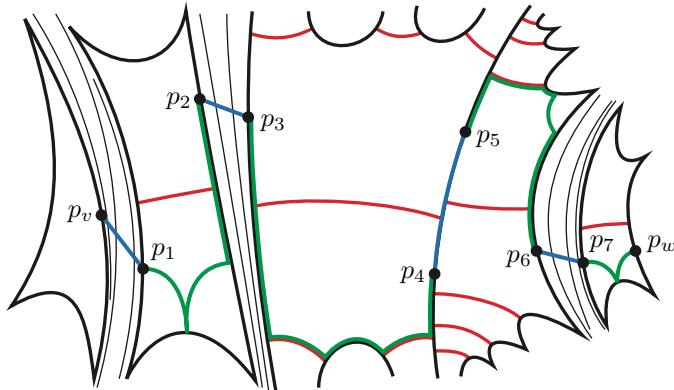


FIGURE 28. Thick subsurfaces between v and w and an admissible route from p_v to p_w . In the figure we have highlighted a path from p_v to p_w through the p_i ; each subpath from p_i to p_{i+1} specifies a factor in the shape-shifting transformation.

We now define an *admissible route* from p_v to p_w to be any sequence of basepoints

$$p_v = p_0, p_1 \in h_1, \dots, p_n \in h_n, p_{n+1} = p_w$$

coming from the projections of the central vertices u_i of hexagons H_{u_i} to $h_i \in \partial_{\lambda} H_{u_i}$. If any geodesic $h_i = h_{i+1}$ is repeated then we require that v and u_i lie on one side of h_i and that w and u_{i+1} lie on the other. Observe that the sequence of pairs (u_i, u_{i+1}) necessarily alternates between simple pairs/pairs sharing a boundary geodesic and pairs which lie in the same (thick) subsurface.

Shape-shifting along admissible routes. To any admissible route we can then define a shape-shifting transformation by concatenating the shape-shifting transformations for subsequent pairs:

$$(43) \quad \varphi_{p_v, p_w} := \varphi_{p_0, p_1} \circ \dots \circ \varphi_{p_n, p_{n+1}}$$

²²Note that this sequence is necessarily finite, as the distance that any geodesic travels in a thick subsurface is bounded below by the shortest arc of $\underline{\alpha}$ (compare the discussion of “close enough” pairs of hexagons in Section 13.2).

where $\varphi_{p_i, p_{i+1}}$ is as in Definition 14.11 if (u_i, u_{i+1}) is simple and as in Definition 14.21 if u_i and u_{i+1} lie in the same subsurface. If $h_i = h_{i+1}$, then we orient h_i to the right as seen from u_i and set $\varphi_{p_i, p_{i+1}} = T_{h_i}^{\mathfrak{s}(u_i, u_{i+1})}$ (recall that we can associate a shear value to the pair (u_i, u_{i+1}) by (15)).

Lemma 14.25. The shape-shifting map φ_{p_v, p_w} is independent of the choice of admissible route (as long as it is defined).

Proof. Since the h_i are uniquely determined, it suffices to change one point at a time.

So suppose that p_i and p'_i are both basepoints on the geodesic h_i ; we then demonstrate the equality

$$\varphi_{p_{i-1}, p_i} \circ \varphi_{p_i, p_{i+1}} = \varphi_{p_{i-1}, p'_i} \circ \varphi_{p'_i, p_{i+1}}$$

from which the lemma follows. Orient h_i so that it runs to the right as seen from v or u_{i-1} .

Without loss of generality, we may assume that the hexagons u_i and u_{i+1} lie in the same subsurface. Otherwise, the hexagons u_{i-1} and u_i lie in the same subsurface and so (p_i, p_{i+1}) is either simple or the points lie on the same isolated leaf. If this happens we prove that

$$\varphi_{p_{i+1}, p_i} \circ \varphi_{p_i, p_{i-1}} = \varphi_{p_{i+1}, p'_i} \circ \varphi_{p'_i, p_{i-1}},$$

which is equivalent to the equation above since each of the shape-shifting factors inverts when one flips the order of the points.

We first consider the shape-shifting transformations coming from comparing p_i or p'_i with p_{i+1} . By our reduction above, u_i and u'_i lie in the same thick subsurface Y . Let $\alpha_1, \dots, \alpha_m$ denote the arcs of $\tilde{\alpha} \cap Y$ encountered when traveling from p'_i to p_i along h_i ; then our definition of shape-shifting in subsurfaces associates the transformation

$$\varphi_{p'_i, p_i} = T_{h_i}^{\varepsilon_1 \sum_{j=1}^m \mathfrak{s}(\alpha_j)}$$

where $\varepsilon_1 = 1$ if h_i is oriented from p'_i to p_i and -1 otherwise. Combining this equation with the subsurface cocycle relation (Proposition 14.24), we have that

$$(44) \quad \varphi_{p'_i, p_{i+1}} = \varphi_{p'_i, p_i} \circ \varphi_{p_i, p_{i+1}} = T_{h_i}^{\varepsilon_1 \sum_{j=1}^m \mathfrak{s}(\alpha_j)} \circ \varphi_{p_i, p_{i+1}}.$$

We now turn our attention to the transformation φ_{p_{i-1}, p'_i} . Consider first the case when (p_{i-1}, p_i) is simple; since p_i and p'_i both lie on h_i this implies that (p_{i-1}, p'_i) is also simple. Moreover, since the geodesics $\mathcal{H}_{i-1, i}$ that separate p_{i-1} from p_i are the same that separate p_{i-1} from p'_i , we may write

$$\varphi_{p_{i-1}, p_i} = \lim_{\mathcal{H} \rightarrow \mathcal{H}_{v, w}} \varphi(s_1) \circ \dots \circ \varphi(s_n) \circ T_{h_i}^{\mathfrak{s}(u_{i-1}, u_i)}$$

and similarly for φ_{p_{i-1}, p'_i} . In particular, each approximation for φ_{p_{i-1}, p_i} differs from the approximation for φ_{p_{i-1}, p'_i} by translation along h_i , and so the same is true in the limit:

$$(45) \quad \varphi_{p_{i-1}, p'_i} = \varphi_{p_{i-1}, p_i} \circ T_{h_i}^{\mathfrak{s}(u_{i-1}, u'_i) - \mathfrak{s}(u_{i-1}, u_i)}.$$

Applying axiom (SH3) for shear-shape cocycles (Definition 7.12) multiple times, we compute that

$$(46) \quad \mathfrak{s}(u_{i-1}, u'_i) - \mathfrak{s}(u_{i-1}, u_i) = \varepsilon_2 \sum_{j=1}^m \mathfrak{s}(\alpha_j)$$

where $\varepsilon_2 = +1$ if p_i precedes p'_i along h_i and -1 if p'_i precedes p_i . Combining (44), (45), and (46),

$$\varphi_{p_{i-1}, p'_i} \circ \varphi_{p'_i, p_{i+1}} = \varphi_{p_{i-1}, p_i} \circ T_{h_i}^{\varepsilon_2 \sum_{j=1}^m \mathfrak{s}(\alpha_j)} \circ T_{h_i}^{\varepsilon_1 \sum_{j=1}^m \mathfrak{s}(\alpha_j)} \circ \varphi_{p_i, p_{i+1}} = \varphi_{p_{i-1}, p_i} \circ \varphi_{p_i, p_{i+1}}$$

since $\varepsilon_2 = -\varepsilon_1$. This completes the proof of the lemma in the case when (u_{i-1}, u_i) is simple.

Similarly, if p_{i-1} and p_i lie on the same isolated leaf of λ then so must p'_i . Unpacking the definitions shows that (45) holds in this case, and Lemma 7.15 implies that (46) does as well. Therefore, in this case we also see that the desired equality holds. \square

Finally, now that we have constructed shape-shifting maps for arbitrary pairs of pointed geodesics in $\partial_\lambda \mathcal{H}$ we can prove that they piece together into an $\text{Isom}^+(\tilde{X})$ -valued cocycle.

Proposition 14.26. The map constructed from X and \mathfrak{s}

$$\begin{aligned}\varphi_{\mathfrak{s}} : \partial_{\lambda}\mathcal{H} \times \partial_{\lambda}\mathcal{H} &\rightarrow \text{Isom}^+(\tilde{X}) \\ ((h_v, p_v), (h_w, p_w)) &\mapsto \varphi_{p_v, p_w}\end{aligned}$$

is a $\pi_1(X)$ -equivariant 1-cocycle, as long as $\|\mathfrak{s}\|_{\tau_{\underline{\alpha}}} < D_{\lambda}(X)/2$.

Proof. That φ is $\pi_1(X)$ -equivariant means that $\varphi_{\gamma p_v, \gamma p_w} = \gamma \circ \varphi_{p_v, p_w} \circ \gamma^{-1}$ for $\gamma \in \pi_1(X)$; this follows directly from the construction.

That φ is a 1-cocycle means it satisfies the familiar cocycle condition on triples, i.e.,

$$\varphi_{p_u, p_v} \circ \varphi_{p_v, p_w} = \varphi_{p_u, p_w}.$$

Note that if p_v lies on some admissible route from p_u to p_w then this is fulfilled automatically by unpacking the definitions and invoking Lemma 14.25.

One special case of the cocycle condition is when $p_u = p_w$; in this case we must show that $\varphi_{p_v, p_w} = \varphi_{p_w, p_v}^{-1}$. To demonstrate this, observe that reversing an admissible route from v to w produces an admissible route from w to v . Moreover, by Corollary 14.14 in the simple case and by definition in the other cases, each $\varphi_{p_i, p_{i+1}}$ also inverts when we flip i and $i+1$, proving that reversing v and w inverts φ_{p_v, p_w} .

Now suppose that u , v , and w are all distinct; then there exists a unique subsurface Y of $\tilde{X} \setminus \tilde{\lambda}$ such that each component of $\tilde{X} \setminus Y$ contains at most one of u , v , or w (note that some of u, v, w may be inside of Y). Choose basepoints r_u , r_v , and r_w on the boundary components of Y that are closest to u , v , and w (if any $\bullet \in \{u, v, w\}$ is in Y then set $r_{\bullet} = p_{\bullet}$). See Figure 29.

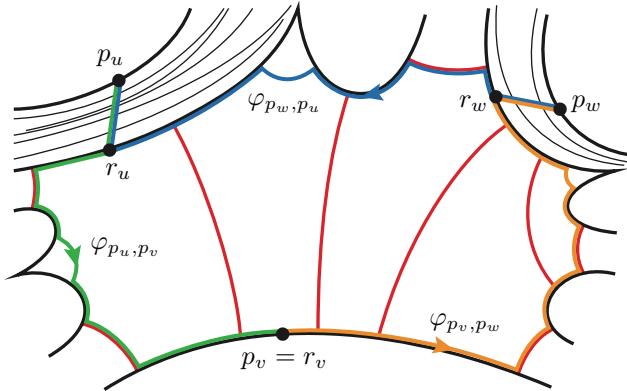


FIGURE 29. The cocycle relation for admissible routes.

Choose an admissible route from p_u to p_v containing r_u and r_v , and similarly for the other two pairs. Then by Lemma 14.25 and the observation that the cocycle condition holds along admissible routes, we may write

$$\varphi_{p_u, p_v} = \varphi_{p_u, r_u} \circ \varphi_{r_u, r_v} \circ \varphi_{r_v, p_v}$$

and similarly for the other two pairs. Combining all three equations and applying the cocycle relation for Y (Proposition 14.24), we see that

$$\begin{aligned}\varphi_{p_u, p_v} \circ \varphi_{p_v, p_w} &= \varphi_{p_u, r_u} \circ \varphi_{r_u, r_v} \circ \varphi_{r_v, p_v} \circ \varphi_{p_v, r_v} \circ \varphi_{r_v, r_w} \circ \varphi_{r_w, p_w} \\ &= \varphi_{p_u, r_u} \circ \varphi_{r_u, r_w} \circ \varphi_{r_w, p_w} = \varphi_{p_u, p_w},\end{aligned}$$

finishing the proof. See Figure 29 for a graphical depiction of this argument. \square

15. SHEAR-SHAPE COORDINATES ARE A HOMEOMORPHISM

We now finish the proof of Theorem 12.1 by proving that the map $\sigma_{\lambda} : \mathcal{T}(S) \rightarrow \mathcal{SH}^+(\lambda)$ is open (Theorem 15.1) and proper and thus, by invariance of domain, a homeomorphism.

In Section 15.1, we use the shape-shifting cocycle $\varphi_{\mathfrak{s}}$, built in the previous section, to deform the representation $\rho : \pi_1(S) \rightarrow \text{PSL}_2 \mathbb{R}$ that induces the hyperbolic structure X . The deformed representation $\rho_{\mathfrak{s}}$ is then

discrete and faithful (Lemma 15.3) and the quotient surface X_s has the desired shear-shape cocycle (Lemma 15.6). In particular, this gives us a continuous local inverse to σ_λ , proving openness. These statements are similar in spirit to those in [Bon96], but the specifics of our proofs are different. In particular, instead of adjusting the relative placements of ideal triangles of $\tilde{X} \setminus \tilde{\lambda}$ we adjust the relative position of pointed geodesics comprising $\tilde{\lambda}$.

We then prove properness of σ_λ in Section 15.2 concluding the proof of Theorem 12.1. Here we return to Bonahon's argument [Bon96, Theorem 20], but applying this strategy in our setting still requires a bit of extra care due to the polyhedral structure of $\mathcal{SH}^+(\lambda)$.

Finally, in Section 15.3 we show that the action of $\mathbb{R}_{>0}$ on $\mathcal{SH}^+(\lambda)$ by dilation produces lines in $\mathcal{T}(S)$ that can sometimes be identified with directed Thurston geodesics.

15.1. Deforming by shape-shifting. In this section, we show that any positive shear-shape cocycle close enough to $\sigma_\lambda(X)$ is actually the geometric shear-shape cocycle of a hyperbolic structure. Compare with [Bon96, Proposition 13].

Theorem 15.1. Let $\underline{\beta}$ be a maximal arc system containing $\underline{\alpha}(X)$ and let $\tau_{\underline{\beta}}$ be a standard smoothing. Then for any $\mathfrak{s} \in W(\tau_{\underline{\beta}})$ such that $\|\mathfrak{s}\|_{\tau_{\underline{\beta}}} < D_\lambda(X)/2$ and such that $\sigma_\lambda(X) + \mathfrak{s}$ represents a positive shear-shape cocycle, there exists $X_s \in \mathcal{T}(S)$ close to X with

$$\sigma_\lambda(X_s) = \sigma_\lambda(X) + \mathfrak{s}.$$

In particular, the image of $\sigma_\lambda(X)$ is open in $\mathcal{SH}^+(\lambda)$.

The proof of this theorem appears at the end of this subsection as the culmination of a series of structural lemmas. Our strategy is to explicitly define X_s by using the shape-shifting cocycle constructed in Section 14 to deform the hyperbolic structure on X . Before proceeding we note the following

Corollary 15.2. For all $t \in \mathbb{R}$, and for all $\mu \in \Delta(\lambda)$, we have the following identity

$$\sigma_\lambda(\text{Eq}_{t\mu}(X)) = \sigma_\lambda(X) + t\mu.$$

Proof. That the earthquake $\text{Eq}_{t\mu}(X)$ is defined for all time is a consequence of countable additivity (equivalently, positivity) of μ ; a complete proof can be found in [EM06, Section III]. Viewing the set of measures supported on λ as a subset of $\mathcal{H}(\lambda)$, the formula is immediate from Theorem 15.1 once we note that $\text{Eq}_{t\mu}(X) = X_{t\mu}$, which follows from the description of $\varphi_{t\mu}$ as a limit of simple left (or right) earthquakes; see (37) and Lemma 14.13. \square

Fix \mathfrak{s} as in the statement of the theorem and pick an arbitrary $v \in \mathcal{H}$ and $(h_v, p_v) \in \partial_\lambda H_v$. Identifying \tilde{X} isometrically with \mathbb{H}^2 and (h_v, p_v) with a pointed line picks out a representation $\rho : \pi_1(S) \rightarrow \text{PSL}_2 \mathbb{R}$ that induces X . Since $\|\mathfrak{s}\|_{\tau_{\underline{\alpha}}} < D_\lambda(X)/2$, Proposition 14.26 allows us to construct the shape-shifting cocycle φ_s .

We may now deform the representation ρ by φ_s by defining

$$\begin{aligned} \rho_s : \pi_1(S) &\rightarrow \text{PSL}_2 \mathbb{R} \\ \gamma &\mapsto \varphi_{p_v, \gamma p_v} \circ \rho(\gamma). \end{aligned}$$

The equivariance and cocycle properties of Proposition 14.26 ensure that ρ_s is itself a representation. Indeed:

$$\begin{aligned} \rho_s(\gamma_1 \gamma_2) &= \varphi_{p_v, \gamma_1 \gamma_2 p_v} \circ \rho(\gamma_1 \gamma_2) \\ &= \varphi_{p_v, \gamma_1 p_v} \circ \varphi_{\gamma_1 p_v, \gamma_1 \gamma_2 p_v} \circ \rho(\gamma_1) \circ \rho(\gamma_2) \\ &= \varphi_{p_v, \gamma_1 p_v} \circ \rho(\gamma_1) \circ \varphi_{p_v, \gamma_2 p_v} \circ \rho(\gamma_1)^{-1} \circ \rho(\gamma_1) \circ \rho(\gamma_2) \\ &= \rho_s(\gamma_1) \circ \rho_s(\gamma_2) \end{aligned}$$

for all $\gamma_1, \gamma_2 \in \pi_1(S)$. Our goal in the rest of the section is then to show that ρ_s is discrete and faithful, and that the quotient surface has the correct geometric shear-shape cocycle.

Adjusting geodesics. To show that ρ_s has the desired properties, we use φ_s to adjust the position of $\tilde{\lambda}$ in \tilde{X} . Ultimately, these adjusted geodesics correspond to the realization of λ on the quotient surface $\tilde{X} / \text{im } \rho_s$.

Let $\mathcal{G}(\tilde{X})$ be the space of geodesics in \tilde{X} , and let $\partial \tilde{\lambda} \subset \mathcal{G}(\tilde{X})$ denote the set of boundary leaves of $\tilde{\lambda}$. Define a map

$$\Phi_{p_v} : \partial \tilde{\lambda} \rightarrow \mathcal{G}(\tilde{X}),$$

as follows: if h is a leaf of $\partial\tilde{\lambda}$, then $h = h_u$ for some (h_u, p_u) in $\partial_\lambda\mathcal{H}$. The map Φ_{p_v} then takes (h_u, p_u) isometrically to the pointed geodesic $\varphi_{p_v, p_u}(h_u, p_u) \subset \tilde{X}$. Note that if $h_u = h_w$ for some other (h_w, p_w) in $\partial_\lambda\mathcal{H}$, then

$$\varphi_{p_v, p_u}^{-1} \circ \varphi_{p_v, p_w} = \varphi_{p_u, p_w}$$

by the cocycle relation (Proposition 14.26) and φ_{p_u, p_w} is by definition a translation along h . Therefore $\Phi_{p_v}h_w = \Phi_{p_v}h_u$, so Φ_{p_v} is indeed well defined.

Using the fact that S is closed, the following lemma follows directly from the fact that ρ_s defines a representation of $\pi_1(S)$ in the same component of representations as ρ . We give a hands on explanation that does not use this fact.

Lemma 15.3. The representation ρ_s constructed above is discrete and faithful.

Proof. For distinct leaves h_u and $h_w \in \partial\tilde{\lambda}$, we claim that $\Phi_{p_v}(h_u)$ is disjoint from $\Phi_{p_v}(h_w)$. Indeed, by the cocycle relation, the position of $\Phi_{p_v}(h_w)$ relative to $\Phi_{p_v}(h_u)$ is the same as the position of $\varphi_{p_u, p_w}h_w$ relative to h_u . Every finite approximation of φ_{p_u, p_w} by compositions of elementary shape-shifting transformations preserves the property that the image of h_w is disjoint from h_u , so the same is true in the limit.

Therefore, as long as $\rho(\gamma)$ does not stabilize h_v then $\Phi_{p_v}(\rho(\gamma)h_v) = \rho_s(\gamma)h_v$ is different from $\Phi_{p_v}(h_v) = h_v$. If $\rho(\gamma)$ is a translation along h_v , we can find γ_0 such that $\rho(\gamma_0\gamma\gamma_0^{-1})h_v \neq h_v$, and so we see $\rho_s(\gamma)$ does not stabilize $\rho(\gamma_0\gamma\gamma_0^{-1})h_v$. In either case, this implies that $\rho_s(\gamma)$ acts nontrivially on the space of geodesics, so in particular $\rho_s(\gamma) \neq 1$, i.e., ρ_s is faithful.

Since $\pi_1(S)$ is a non-elementary group and ρ_s is faithful, $\text{im } \rho_s$ is a non-elementary subgroup of isometries. So assume towards contradiction that ρ_s is indiscrete. Then $\text{im } \rho_s$ must be dense in $\text{PSL}_2 \mathbb{R}$; see, e.g., [Sul85, Proposition, p. 246]. In particular, there is an element $\gamma \in \pi_1(S)$ such that $\rho_s(\gamma)$ is arbitrarily close to a rotation of angle $\pi/2$ around p_v . Then $\rho_s(\gamma)h_v = \Phi_{p_v}\rho(\gamma)h_v$ meets h_v in a point, which is impossible, because $\Phi_{p_v}(h_v)$ is either equal to or disjoint from $\Phi_{p_v}(\rho(\gamma)h_v)$. We conclude that ρ_s is discrete, completing the proof of the lemma. \square

By Lemma 15.3, the quotient $X_s = \tilde{X}/\text{im } \rho_s$ is a hyperbolic surface equipped with a homeomorphism $S \rightarrow X_s$ in the homotopy class determined by ρ_s . As such, ρ_s induces a (ρ, ρ_s) -equivariant homeomorphism $\partial\tilde{X} \rightarrow \partial\tilde{X}$, hence a continuous, equivariant map on the space of geodesics.

Lemma 15.4. The map Φ_{p_v} extends continuously to $\tilde{\lambda}$, and $\Phi_{p_v}(\tilde{\lambda})$ descends to the geodesic realization of λ on X_s .

Proof. By equivariance, the induced map on geodesics agrees with Φ_{p_v} on $\partial\tilde{\lambda}$. The leaves of $\partial\tilde{\lambda}$ are dense in $\tilde{\lambda}$, so the closure of the image of Φ_{p_v} is the geodesic realization of λ on \tilde{X}_s , which is invariant under the action of ρ_s . \square

Since $\Phi_{p_v}(\tilde{\lambda})$ is the lift of the realization of λ on X_s , we may now leverage our understanding of the shape-shifting cocycle to show that the complementary subsurfaces of $X_s \setminus \lambda$ have the desired shapes.

Lemma 15.5. We have $\underline{A}(X_s) = \underline{A}(X) + \mathfrak{a}$.

Proof. Recall that by construction the unweighted arc systems of $X \setminus \lambda$ and $X_s \setminus \lambda$ are both contained in some joint maximal arc system β , leading to an identification of hexagons of $\tilde{X} \setminus \lambda$ with those of $\tilde{X} \setminus \lambda_s$.

So let β be an arc of β , realized orthogeodesically in X_s . Let β also denote a choice of lift, orthogonal to λ_s in \tilde{X} , and let u and w denote the hexagons adjacent to β . Choose either of the geodesics g of λ meeting β and let q_u and q_w be the basepoints of u and w on g . Then by the cocycle relation (Proposition 14.26), we know that

$$\varphi_{p_v, p_w} = \varphi_{p_v, p_u} \circ \varphi_{p_u, p_w}$$

and so applying equivariance we see that the placement of q_w relative to q_u differs from the placement of p_w relative to p_u only by φ_{p_u, p_w} .

But now since u and w are in the same subsurface, we see by definition that φ_{p_u, p_w} is translation along g by exactly $\mathfrak{s}(\beta)$. Therefore, the distance along $\varphi_{p_u, p_w}g$ between q_u and q_w is exactly the distance along g between p_u and p_w plus $\mathfrak{s}(\beta)$. Translated into arc weights,

$$\sigma_\lambda(X_s)(\beta) = \sigma_\lambda(X)(\beta) + \mathfrak{s}(\beta),$$

completing the proof of the lemma. \square

Now that we know that the “shape” part of the data of $\sigma_\lambda(X_s)$ is what it is supposed to be, we need only check that the “shearing” data is as specified. Compare [Bon96, Lemma 19].

Lemma 15.6. The surface X_s has geometric shear-shape cocycle $\sigma_\lambda(X_s) = \sigma_\lambda(X) + \mathfrak{s}$.

Proof. Observe that by the cocycle relation (Proposition 14.26) and the discussion in Section 13.2 it suffices to compute the change in shearing data between simple pairs.

So suppose that (v, w) is simple. For each integer r , recall that $\mathcal{H}_{v,w}^r = (u_i)_{i=1}^n$ denotes the set of hexagons such that the intersection of the geodesic from p_v to p_w with u_i has depth at most r with respect to a fixed geometric train track. Set $v = u_0$ and $w = u_{n+1}$, and let $h_i = g_{u_i}^v$, the pointed boundary geodesic of u_i closest to v . Then by Lemma 14.13, we know that

$$\varphi_{p_v, p_w}^r = T_{h_0}^{\mathfrak{s}(u_0, u_1)} \circ A(s_1) \circ T_{h_1}^{\mathfrak{s}(u_1, u_2)} \circ A(s_2) \circ \dots \circ A(s_n) \circ T_{h_n}^{\mathfrak{s}(u_n, u_{n+1})}$$

is a good approximation of φ_{p_v, p_w} for large enough r .

Now for each r we can deform the hyperbolic structure on \tilde{X} by φ_{p_v, p_w}^r (sacrificing equivariance) and measure the shear $\sigma_r(v, w)$ between v and w in that deformed structure. More precisely, we recall that if h'_i denotes the other geodesic in u_i that separates v from w , the spike-shaping transformation is equal to a translation along h'_i then along h_i . We may then deform \tilde{X} by replacing each translation in the factorization of φ_{p_v, p_w}^r with a (right) earthquake along the same geodesic; compare with our “geometric explanation” of spike-shaping in Section 14.2.

Since each translation $T_{h_i}^{\mathfrak{s}(u_i, u_{i+1})}$ appearing in φ_{p_v, p_w}^r shears \tilde{X} along a leaf of λ , it preserves the orthogeodesic foliation in complementary components. Therefore, each such term in the deformation thus changes the shear between v and w by exactly $\mathfrak{s}(u_i, u_{i+1})$.

On the other hand, each spike-shaping transformation $A(s_i)$ is a parabolic transformation fixing the vertex of the spike and thus preserves horocycles based at that point. In particular, the distinguished basepoints of each h_i and h'_i remain on the same horocycle and hence deforming by $A(s_i)$ does not affect $\sigma_r(v, w)$.

In summary, deforming \tilde{X} by the approximation φ_{p_v, p_w}^r changes the shear between v and w by

$$\sigma_r(v, w) - \sigma(v, w) = \sum_{i=0}^n \mathfrak{s}(u_i, u_{i+1}) = \mathfrak{s}(v, w)$$

where the last equality follows from finite additivity (axiom (SH2)).

Since this equality holds in each approximation and $\varphi_{p_v, p_w}^r \rightarrow \varphi_{p_v, p_w}$ as $r \rightarrow \infty$, the equality holds in the limit as well. Therefore, deforming \tilde{X} by φ_{p_v, p_w} changes the shear between v and w by exactly $\mathfrak{s}(v, w)$, which is what we needed to show. \square

Proof of Theorem 15.1. As $\|\mathfrak{s}\|_{\tau_\alpha} < D_\lambda(X)/2$, Lemmas 14.10 and 14.13 ensure that the limits in the definition of φ_{p_v, p_w} make sense for all simple pairs (v, w) . Proposition 14.26 then allows us to construct φ_s . By Lemma 15.3, the deformed representation $\rho_s = \varphi_s \cdot \rho$ is discrete and faithful, and by Lemma 15.6 the quotient surface has the correct geometric shear-shape cocycle.

Finally, we observe that the values of shape-shifting cocycle φ_s all converge to the identity as $\|\mathfrak{s}\|_{\tau_\alpha} \rightarrow 0$, and consequently $X_s \rightarrow X$. This completes the proof of the theorem. \square

15.2. The global structure of the shear-shape map. We have already proven in Proposition 13.14 that the image of σ_λ lies in $\mathcal{SH}^+(\lambda)$. We now show that this containment is in fact an equality, completing the proof of Theorem 12.1.

We proceed in two steps; the first is to show that

Proposition 15.7. The shear-shape map σ_λ is a homeomorphism onto its image.

Proof. Proposition 13.12 (injectivity of σ_λ) allows us to invert σ_λ on its image and for each $X \in \mathcal{T}(S)$, Theorem 15.1 provides us with an open neighborhood of $\sigma_\lambda(X) \in \mathcal{SH}^+(\lambda)$ on which σ_λ^{-1} is defined and continuous. By Proposition 8.5, we have that $\mathcal{SH}^+(\lambda) \subset \mathcal{SH}(\lambda)$ is an open cell of dimension $6g - 6$. Invoking invariance of domain, we get that σ_λ^{-1} and hence σ_λ are local homeomorphisms. An additional application of Proposition 13.12 implies that σ_λ is globally injective, so σ_λ is a homeomorphism onto its image as claimed. \square

The second step is to prove that $\sigma_\lambda : \mathcal{T}(S) \rightarrow \mathcal{SH}^+(\lambda)$ is a proper map. That is, we must show that when X_k escapes to infinity in $\mathcal{T}(S)$, the corresponding shear-shape cocycles $\sigma_\lambda(X_k)$ must diverge in $\mathcal{SH}^+(\lambda)$. Since proper local homeomorphisms are coverings and $\mathcal{SH}^+(\lambda)$ is a cell, the map σ_λ must be a homeomorphism.

The proof we present below is essentially just that of [Bon96, Theorem 20], but we have to address the additional complications introduced by the PL structure of $\mathcal{SH}^+(\lambda)$; this manifests itself in the *stratified* real-analytic structure of the map. ²³

Proof of Theorem 12.1. We begin by recording an estimate for the geometry of surfaces near the boundary of the image of σ_λ (where “near” is measured in a train track chart).

So suppose that $X \in \mathcal{T}(S)$, set $\underline{\alpha} = \underline{\alpha}(X)$, and build a standard smoothing $\tau_{\underline{\alpha}}$ carrying λ geometrically on X . Fix $\epsilon > 0$ and suppose that there exists some $\mathfrak{s} \in W(\tau_{\underline{\alpha}})$ with $\|\mathfrak{s}\|_{\tau_{\underline{\alpha}}} < \epsilon$ such that $\sigma_\lambda(X) + \mathfrak{s} \in \mathcal{SH}^+(\lambda)$ is not in the image of σ_λ ; then Theorem 15.1 implies that

$$D_\lambda(X)/2 \leq \|\mathfrak{s}\|_{\tau_{\underline{\alpha}}} < \epsilon.$$

The following claim can be extracted from the proof of [Bon96, Theorem 20]; we outline a proof for the convenience of the reader.

Claim 15.8. There is a transverse measure $\mu \in \Delta(\lambda)$ with $\frac{1}{9\chi(S)} \leq \|\mu\|_{\tau_{\underline{\alpha}}} \leq 1$ and

$$\ell_X(\mu) = \omega_{\mathcal{SH}}(\sigma_\lambda(X), \mu) < \epsilon.$$

Proof of Claim 15.8. If there is a simple closed curve component of λ with length at most ϵ , then we are done. Otherwise, even though $\underline{A} + \mathfrak{a}$ defines a hyperbolic structure on each piece of $S \setminus \lambda$, the overall shear-shape cocycle $\sigma_\lambda(X) + \mathfrak{s}$ does not define a hyperbolic structure on S because the proof of Lemma 14.10 or Lemma 14.13 fails. Therefore, there is a simple pair (v, w) for which the finite products $\varphi_{\mathcal{H}}$ (or $\varphi_{v,w}^r$) fail to converge as \mathcal{H} tends to $\mathcal{H}_{v,w}$ (or $r \rightarrow \infty$).

We claim that there exists u between v and w and a spike $s = (g, h)$ of H_u such that the following holds: for any geodesic transversal $k \subset X$ to λ meeting the spike s , the countably many points of $\tilde{k} \cap g \subset g$ (labeled by $r \in \mathbb{N}$) exiting one end of g escape at a rate strictly slower than $\epsilon(r-1)$. In other words, there are segments $d_r \subset g$ such that $\ell_X(d_r) \leq \epsilon(r-1)$ and d_r meets k exactly r times.

If this were not the case, then as in the proof of Lemma 14.5, the “gaps” $c_r \subset k_{v,w} \setminus \lambda$ have length $\ell_X(c_r) = O(e^{-\epsilon r})$, where $c_r \cap g$ is labelled by $r \in \mathbb{N}$. This estimate on the decay of gaps implies that $\varphi_{\mathcal{H}}$ converges as $\mathcal{H} \rightarrow \mathcal{H}_{v,w}$ and that $\varphi_{v,w}^r \rightarrow \varphi_{v,w}$ as $r \rightarrow \infty$ (see the proof of Lemma 14.10), contradicting our assumption.

Now consider the weight system w_r on $\tau_{\underline{\alpha}}$ (not satisfying the switch conditions) defined by counting the number of times d_r travels along each branch of $\tau_{\underline{\alpha}}$, and dividing by the total number of branches n_r that d_r traverses, with multiplicity. Observe that $n_r \geq r$ by definition. Then $\|w_r\|_{\tau_{\underline{\alpha}}} \leq 1$ in the vector space $\mathbb{R}^{b(\tau_{\underline{\alpha}})}$ and w_r takes value zero on branches corresponding to arcs of $\underline{\alpha}$. Moreover, w_r is non-negative on each branch and approaches the weight space $W(\tau) \subset \mathbb{R}^{b(\tau)}$ as $r \rightarrow \infty$. Since w_r are built from leaves of λ , any limit point μ defines a transverse measure supported on λ (compare also [PH92, Proposition 3.3.2]).

There are at most $9\chi(S)$ branches of $\tau_{\underline{\alpha}}$, so by the pigeonhole principle there is a branch such that each w_r has mass at least $1/9\chi(S)$, and therefore so must μ . But now by construction,

$$\ell_X(\mu) = \lim_{r \rightarrow \infty} \frac{\ell_X(d_r)}{n_r} < \frac{(r-1)\epsilon}{n_r} < \epsilon,$$

providing the desired measure. \square

Now suppose towards contradiction that $\underline{\alpha}$ is maximal and $\sigma_\lambda(X_k) \in \mathcal{SH}^+(\lambda; \underline{\alpha})$ is a sequence approaching some $\sigma \in \mathcal{SH}^+(\lambda; \underline{\alpha})$ that is not in the image of σ_λ . We may then apply the above construction to $\sigma - \sigma_\lambda(X_k)$ to extract a family of measures μ_k on λ satisfying $1/9\chi(S) \leq \|\mu_k\|_{\tau_{\underline{\alpha}}} \leq 1$ and $\omega_{\mathcal{SH}}(\sigma_\lambda(X), \mu_k) \rightarrow 0$. By compactness of the set measures on λ with norm bounded away from zero and infinity, there is some non-zero accumulation point μ of μ_k . Continuity of $\omega_{\mathcal{SH}}$ (Lemma 8.3) then gives

$$\omega_{\mathcal{SH}}(\sigma_\lambda(X_k), \mu_k) \rightarrow \omega_{\mathcal{SH}}(\sigma, \mu) = 0,$$

²³Recall that $\mathcal{H}^+(\lambda)$ is an open cone with finitely many faces in a vector space, while $\mathcal{SH}^+(\lambda)$ is an affine cone bundle over a piecewise linear space with no obvious way of extending the smooth structure over faces of $\mathcal{B}(S \setminus \lambda)$.

and so we see that $\sigma \notin \mathcal{SH}^+(\lambda; \underline{\alpha})$, a contradiction. Hence $\text{im}(\sigma_\lambda) \cap \mathcal{SH}^+(\lambda; \underline{\alpha})$ is relatively closed. On the other hand, σ_λ is a local homeomorphism by Proposition 15.7, hence $\text{im}(\sigma_\lambda) \cap \mathcal{SH}^+(\lambda; \underline{\alpha})$ is relatively open.

If we knew that the projection of $\text{im}(\sigma_\lambda)$ surjects onto $\mathcal{B}(S \setminus \lambda)$ (or at least meets each top-dimensional face) we would be done. Since we do not *a priori* have this information, we instead work our way out in $\mathcal{B}(S \setminus \lambda)$ cell by cell.

To wit, we may invoke Theorem 15.1 once more to deduce that $\text{im}(\sigma_\lambda) \cap \mathcal{SH}^+(\lambda; \underline{\alpha}')$ is relatively open for every filling arc system $\underline{\alpha}'$ that shares a common filling arc subsystem with $\underline{\alpha}$ (hence $\mathcal{SH}^+(\lambda; \underline{\alpha})$ and $\mathcal{SH}^+(\lambda; \underline{\alpha}')$ intersect). Repeating the argument above for these cells, we have that $\text{im}(\sigma_\lambda) \supset \mathcal{SH}^+(\lambda; \underline{\alpha}')$ as well. Since $\mathcal{B}(S \setminus \lambda)$ is connected, iterating this procedure allows us to deduce that $\text{im}(\sigma_\lambda) \supset \mathcal{SH}^+(\lambda)$. The reverse inclusion follows from Corollary 13.14, so σ_λ is a homeomorphism onto $\mathcal{SH}^+(\lambda)$.

To address the regularity of σ_λ , we note that while $\mathcal{T}(S)$ has a natural \mathbb{R} -analytic structure, $\mathcal{SH}(\lambda)$ does not. However, for each arc system $\underline{\alpha}$, filling or not, the open cell $\mathcal{B}^\circ(\underline{\alpha})$ has a well defined analytic structure compatible with that of the analytic submanifold of $\mathcal{T}(S \setminus \lambda)$ that it parametrizes. The total space of the bundle $\mathcal{SH}^\circ(\lambda; \underline{\alpha}) \rightarrow \mathcal{B}^\circ(\underline{\alpha})$ also carries an analytic structure, invariant under train track coordinate-transformations (Proposition 8.5); thus $\mathcal{SH}(\lambda)$ has a stratified \mathbb{R} -analytic structure.

The shape-shifting cocycle φ_s , hence the surface X_s , then depends real-analytically on $s \in W(\tau_\underline{\alpha})$ (where $\underline{\alpha}$ here is equal to the support of $\underline{A}(X)$, not a maximal completion). The reason for this is clear: all elementary shape-shifting transformations are products of small parabolic transformations (see [Thu98, Section 9] or [Bon96, Theorem A]) or translations with translation distance that are (restrictions of) real-analytic functions on (an analytic submanifold of) $\mathcal{T}(S \setminus \lambda)$. These products converge absolutely to the shape-shifting cocycle, hence uniformly on compact sets to an analytic deformation. \square

15.3. Dilation rays and Thurston geodesics. Using our coordinatization, we can define an extension of the earthquake flow to an action by the upper-triangular subgroup.

Definition 15.9. Given a measured geodesic lamination λ , a hyperbolic surface $X \in \mathcal{T}(S)$, and $t \in \mathbb{R}$, define an analytic path of surfaces $\{X_\lambda^t\}_{t \in \mathbb{R}}$ by

$$X_\lambda^t := \sigma_\lambda^{-1}(e^t \sigma_\lambda(X)),$$

called the *dilation ray*²⁴ based at X directed by λ .

As the earthquake flow acts by translation in coordinates (Corollary 15.2), we see that dilation and earthquake along λ (together with scaling the measure on λ) fit together into an action by the upper-triangular subgroup $B < \text{GL}_2^+ \mathbb{R}$ on \mathcal{PT}_g . More explicitly, we can specify an action of B on $\mathcal{T}(S) \times \mathbb{R}_{>0} \lambda$ (by homeomorphisms) by setting

$$(47) \quad \begin{pmatrix} a & b \\ 0 & c \end{pmatrix} \cdot (X, \lambda) := (\sigma_\lambda^{-1}(a\sigma_\lambda(X) + b\lambda), c\lambda).$$

These B -actions assemble into a $\text{Mod}(S)$ -equivariant B -action on \mathcal{PT}_g (observe that σ_λ depends only on the support of λ and not the actual measure). Quotienting by the mapping class group and restricting to the unit length locus then gives a P -action on $\mathcal{P}^1 \mathcal{M}_g$, and since dilation preserves the property of being regular, a P -action on each stratum $\mathcal{P}^1 \mathcal{M}_g^{\text{reg}}(\underline{\kappa})$. We call any such action an action by *stretchquakes*.

Using the commutativity of Diagram (2) (Theorem 13.13), we can compare (47) with the computations performed in Lemmas 11.1 and 11.2 to see that

Proposition 15.10. The map \mathcal{O} takes the P action of (47) on $\mathcal{P}^1 \mathcal{M}_g$ to the standard P action on $\mathcal{Q}^1 \mathcal{M}_g$.

While we have defined them via coordinates, it is not hard to see that dilation rays are geometrically meaningful families of surfaces. Generally, we obtain paths of surfaces along which the length of λ scales nicely, and we can identify some dilation rays as directed lines in Thurston's asymmetric metric on $\mathcal{T}(S)$.

Mirzakhani observed (see [Mir08, Remark p. 33]) that for a maximal lamination μ , the dilation ray $t \mapsto X_\mu^t$ corresponds to the stretch path directed by μ defined by Thurston in [Thu98, Section 4]. Very roughly, stretch paths are obtained by gluing together certain expanding self-homeomorphisms of the ideal triangles comprising $X \setminus \mu$ along the leaves of μ .

²⁴We are abusing terminology here by declaring that the image of \mathbb{R} under an analytic mapping is a ray. Our aim is to emphasize that the dilation ray should be thought of as directed toward the future, even though it can be defined for all time.

Lemma 15.11 (Proposition 2.2 of [Thu98]). Let P_n be a regular ideal hyperbolic n -gon. For any $K \geq 1$, there is a K -Lipschitz self-homeomorphism $P_n \rightarrow P_n$ that maps each side to itself and expands arclength along the boundary by a constant factor of K .

Proof. The orthogeodesic foliation $\mathcal{O}(P_n)$ is measure equivalent to a partial foliation by horocycles centered at the spikes of P_n . The desired K -Lipschitz homeomorphism $P_n \rightarrow P_n$ is constructed by fixing the central horocyclic n -gon and mapping each horocyclic arc at distance s from the central region to the horocyclic arc at distance Rs in the same spike. \square

Any partition $\underline{\kappa} = (\kappa_1, \dots, \kappa_n)$ of $4g - 4$ determines a regular locus $\mathcal{PT}_g^{\text{reg}}(\underline{\kappa})$ of pairs (X, λ) , where the complement of λ in X is a union of regular ideal $(\kappa_i + 2)$ -gons. Then $\mathcal{P}^1 \mathcal{M}_g^{\text{reg}}(\underline{\kappa})$ is the moduli space of pairs where $\ell_X(\lambda) = 1$.

Gluing together the expanding maps of regular polygons provides an explicit model of dilation rays in $\mathcal{P}^1 \mathcal{M}_g(\underline{\kappa})$ and identifies them with geodesics for the Thurston metric. A survey of some basic properties of Thurston's metric as well as similarities and differences between directed stretch rays and Teichmüller geodesics can be found in [PT07]. The following proposition was inspired in part by recent work of Horbez and Tao, in which they investigate the minimally displaced sets in the Thurston's metric using a similar construction [HT].

Proposition 15.12. For any $(X, \lambda) \in \mathcal{PT}_g^{\text{reg}}(\underline{\kappa})$, the dilation ray $\{X_\lambda^t : t \in \mathbb{R}\} \subset \mathcal{PT}_g^{\text{reg}}(\underline{\kappa})$ is a directed unit-speed geodesic in Thurston's asymmetric Lipschitz metric.

Proof. Since λ is regular on X , $\sigma_\lambda(X) \in \mathcal{SH}^+(\lambda)$ lies in the fiber over the empty arc system. Scaling $\sigma_\lambda(X)$ preserves this arc system, so X_λ^t is regular for all t . It suffices to prove that the optimal Lipschitz constant for a map $X \rightarrow X_\lambda^t$ in the homotopy class determined by markings is e^t for all $t \geq 0$.

Let $H_\lambda(X)$ denote the (partial) foliation of X by horocyclic arcs that is measure equivalent to $\mathcal{O}_\lambda(X)$. The maps of Lemma 15.11 assemble to an e^t -Lipschitz homeomorphism $X \setminus \lambda \rightarrow X_\lambda^t \setminus \lambda$ such that $H_\lambda(X)$ maps to $H_\lambda(X_\lambda^t) = e^t H_\lambda(X)$ on each component (as measured foliations). Now, using the fact that $\sigma_\lambda(X_\lambda^t) = e^t \sigma_\lambda(X)$, we can adapt the argument of [Bon96, Lemma 11] (as sketched in Proposition 13.12) to show that this map is locally Lipschitz hence extends across λ to an e^t -Lipschitz homeomorphism $X \rightarrow X_\lambda^t$.

Thus e^t provides an upper bound for the optimal Lipschitz constant in the homotopy class determined by markings. On the other hand,

$$\ell_{X_\lambda^t}(\lambda) = \omega_{\mathcal{SH}}(\sigma_\lambda(X_\lambda^t), \lambda) = \omega_{\mathcal{SH}}(e^t \sigma_\lambda(X), \lambda) = e^t \ell_X(\lambda),$$

so e^t is also a lower bound for the optimal Lipschitz constant. This completes the proof of the proposition. \square

Remark 15.13. As in the last line of the proof of Proposition 15.12 we always have $\ell_{X_\lambda^t}(\lambda) = e^t \ell_X(\lambda)$ for arbitrary $\lambda \in \mathcal{ML}(S)$. Thus the distance from X to X_λ^t in Thurston's metric is at least t . However, we do not always know how to build e^t -Lipschitz proper homotopy equivalences $X \setminus \lambda \rightarrow X_\lambda^t \setminus \lambda$ (in the correct homotopy class) that expand arclength along $\partial X \setminus \lambda$ by a constant factor of e^t .

Remark 15.14. *Added in proof:* In recent work, Pan and Wolf build new families of geodesics for the Lipschitz metric using harmonic maps [PW22]. Their work also uses our coordinates to show that certain “Hopf differential disks” in \mathcal{T}_g converge to “stretch–earthquake disks.” It would be interesting to know if their new geodesics coincide with the dilation rays defined here, and by extension if their stretch–earthquake disks are the same as the orbits of the stretchquake action defined here.

Remark 15.15. Our dilation rays are different from Thurston's stretch rays defined with respect to one of the finitely many maximal completions of λ when λ is not maximal. This follows from the fact that $\mathcal{O}_\lambda(X) \neq \mathcal{O}_{\lambda'}(X)$, where λ' is a maximal completion of λ .

The map $\mathcal{PT}_g^{\text{reg}}(\underline{\kappa}) \times \mathbb{R} \rightarrow \mathcal{PT}_g^{\text{reg}}(\underline{\kappa})$ defined by the rule $(X, \lambda, t) \mapsto (X_\lambda^t, e^{-t}\lambda)$ is called the *stretch flow*. The stretch flow is $\text{Mod}(S)$ -equivariant and

$$\ell_{X_\lambda^t}(e^{-t}\lambda) = \ell_X(\lambda),$$

hence descends to $\mathcal{P}^1 \mathcal{M}_g^{\text{reg}}(\underline{\kappa})$.

Corollary 15.16. Let ν be a P -invariant ergodic probability measure on $\mathcal{P}^1 \mathcal{M}_g$.

- For ν -almost every (X, λ) , the dilation ray $t \mapsto X_\lambda^t$ is a unit speed geodesic in Thurston's asymmetric metric.
- On a set of full ν -measure, the action of the diagonal subgroup of P is identified with the stretch flow and \mathcal{O} conjugates stretch flow to Teichmüller geodesic flow.

In particular, the stretch flow is ergodic with respect to ν .

Proof. By Corollary 1.2, ν -almost every point is regular (with respect to the same topological type of lamination), so the first statement of the theorem is immediate from Proposition 15.12.

The second statement is essentially a restatement of Theorem B combined with the previous statement. Alternatively, in the Gardiner-Masur parameterization of \mathcal{QT}_g (Theorem 2.1), the Teichmüller geodesic flow at time t is given by $(\eta, \lambda) \mapsto (e^t \eta, e^{-t} \lambda)$, so unraveling the definitions and using commutativity of Diagram (2) (Theorem 13.13) gives the result.

For ergodicity, we apply Theorem C which asserts, in particular, that $\mathcal{O}_* \nu$ is an ergodic $\text{SL}_2 \mathbb{R}$ -invariant probability measure on $\mathcal{QM}_g(\kappa)$. The Howe-Moore Theorem implies that any non-compact, closed subgroup of $\text{SL}_2 \mathbb{R}$ inherits ergodicity (see, e.g., [FK02, Theorem 3.3.1]); in particular, the Teichmüller geodesic flow is ergodic with respect to $\mathcal{O}_* \nu$. So \mathcal{O} maps any stretch flow-invariant set B of positive ν -measure to an $\mathcal{O}_* \nu$ Teichmüller geodesic flow-invariant set of positive measure, which must have full measure by ergodicity. Thus $\nu(B) = 1$, demonstrating ergodicity of the stretch flow. \square

Recently, Allessandrini and Disarlo [AD20] constructed Lipschitz maps between some pairs of degenerate right angled hexagons that stretch alternating boundary geodesics by a constant factor. Recall from Section 6 that the Teichmüller space of an ideal quadrilateral is 1 dimensional and can be described as the the cone over a pair of points corresponding to the two arcs α and β that join opposite sides of Q .

Lemma 15.17. Let Q be an ideal quadrilateral with weighted filling arc system $s\delta$, where $\delta \in \{\alpha, \beta\}$. Let Q^t be the quadrilateral with arc system $e^t s\delta$. There is an e^t -Lipschitz surjection $Q \rightarrow Q^t$ that multiplies arclength along the boundary of Q by a factor of e^t . Moreover, the projection of the compact edge of the spine of Q is mapped to the projection of the compact edge of the spine of Q^t .

Proof. Every ideal quadrilateral has an orientation preserving isometric involution swapping opposite sides. Thus the orthogeodesic representative of δ cuts Q into 2 isometric pieces, each of which is a right angled hexagon with two degenerate sides. On each piece, we can apply [AD20, Lemma 6.9] to obtain maps which glue together along δ to give a map with the desired properties. \square

We immediately obtain some new geodesics for Thurston's metric.

Proposition 15.18. If $S \setminus \lambda$ consists of ideal triangles and quadrilaterals, then for any $X \in \mathcal{T}(S)$, $t \mapsto X_\lambda^t$ is a directed, unit speed geodesic for Thurston's asymmetric metric.

Proof. The proof is nearly identical to the proof of Proposition 15.12 so we only provide a brief outline. Construct an e^t -Lipschitz surjective map $X \setminus \lambda \rightarrow X_\lambda^t \setminus \lambda$ from the units of Lemma 15.11 and Lemma 15.17. For the same reason as before, this map extends continuously across the leaves of λ and provides an e^t -Lipschitz homotopy equivalence $X \rightarrow X_\lambda^t$ in the homotopy class determined by markings. Thus e^t is an upper bound for the Lipschitz constant among homotopy equivalences $X \rightarrow X_\lambda^t$ in correct homotopy class. This is clearly an upper bound for the ratio

$$\max_{\mu \in \mathcal{ML}(S)} \frac{\ell_\mu(X_\lambda^t)}{\ell_\mu(X)}.$$

But e^t is also a lower bound for this ratio, because the length of λ is scaled by a factor of e^t .

By a theorem of Thurston [Thu98, Theorem 8.5], there is a e^t -Lipschitz homeomorphism $X \rightarrow X_\lambda^t$ homotopic to the map constructed above. This completes the proof of the proposition. \square

Remark 15.19. The proof of Proposition 15.18 clearly supplies a more general statement: If λ is filling and cuts $X \in \mathcal{T}(S)$ into a regular polygons and quadrilaterals of any shape, then $t \mapsto X_\lambda^t$ is a geodesic for Thurston's metric.

There are other cases in which we can glue Lipschitz maps between degenerate right angled hexagons that can be found in the literature (e.g., [AD20, PY17]). However, these other cases require additional symmetry that is not always present in our setting. We suspect that there is a different approach that would prove

that dilation rays can always be identified with Thurston geodesics, so that \mathcal{O} conjugates a kind of Thurston geodesic flow to Teichmüller geodesic flow.

16. FUTURE AND ONGOING WORK

There is much more to understand about the correspondence between hyperbolic and flat geometry described in this paper. In addition to using the orthogeodesic foliation to import tools from Teichmüller dynamics into the world of hyperbolic geometry (and vice versa), the authors expect this link to provide retroactive explanations for analogous phenomena in the two settings.

We describe a number of future directions and potential applications of the correspondence below, some of which will be addressed in forthcoming sequels.

Continuity and equidistribution. Theorem D states that for a fixed lamination $\mathcal{O}_\lambda : \mathcal{T}(S) \rightarrow \mathcal{MF}(\lambda)$ is a homeomorphism, but as Mirzakhani already observed [Mir08, p. 33], \mathcal{O} cannot be continuous on \mathcal{PT}_g . Moreover, Arana-Herrera and Wright have proven that the earthquake and horocycle flow are not topologically conjugate by any map [AHW22]. At fault is the basic fact that the support of a measured lamination does not vary continuously in the relevant topology.

In forthcoming work [CFa], the authors investigate the continuity properties of \mathcal{O} restricted to specific families of (X, λ) with constrained geometry and topology. On these families, the support of λ is forced to vary continuously in the Hausdorff topology as the pair varies (in the usual topology on \mathcal{PT}_g). For example, each of the regular loci has this property. With this extra geometric control in hand, we prove that \mathcal{O} restricts to a homeomorphism $\mathcal{PT}_g^{\text{reg}}(\underline{\kappa}) \leftrightarrow \mathcal{QT}_g^{\text{nsc}}(\underline{\kappa})$ on each regular locus.

By imposing a stronger (yet still geometrically meaningful) topology on $\mathcal{ML}(S)$, we ensure the continuity of \mathcal{O} varying over all pairs: let $\mathfrak{ML}(S)$ denote the set of measured laminations with the “Hausdorff + measure” topology so that measured laminations are close in $\mathfrak{ML}(S)$ if they are close both in measure and their supports are Hausdorff close. We prove a general phenomenon that $\mathcal{O} : \mathcal{T}(S) \times \mathfrak{ML}(S) \rightarrow \mathcal{QT}(S)$ is locally Hölder continuous with respect to a nice family of locally defined metrics in geometric train track coordinates.

Our continuity arguments depend on a detailed analysis of the geometric structure of small foliated train track neighborhoods of a lamination on a hyperbolic surface. This analysis is sufficiently robust to produce “enough continuity” to deduce that \mathcal{O} is a Borel-measurable isomorphism, a fact which is pivotal for applications. The results of Section 12 then live in a more natural setting, as well.

Combined with this work, the conjugacy of Theorems A and B allows us to import techniques of flat geometry to the hyperbolic setting. In particular, while \mathcal{O} is not continuous, its discontinuity is controlled enough that we can translate between equidistribution in \mathcal{PT}_g and equidistribution in \mathcal{QT}_g [CFb].

Symplectic structure. For a maximal lamination λ , Bonahon and Sözen identified the Goldman symplectic form on the Teichmüller component of $\text{Hom}(\pi_1 S, \text{PSL}_2 \mathbb{R}) / \text{PSL}_2 \mathbb{R}$ (also the Weil-Petersson symplectic form) as $\sigma_\lambda^* \omega_H$ in shear coordinates [SB01]. For arbitrary $\lambda \in \mathcal{ML}(S)$ and $X \in \mathcal{T}(S)$, the shape-shifting cocycles built in Section 14 provide an open set of deformations of the hyperbolization $[\rho : \pi_1 S \rightarrow \text{PSL}_2 \mathbb{R}]$ of X (Theorem 15.1). Taking derivatives (as in [SB01]) identifies the tangent space to $[\rho]$ with the vector space of Ad_ρ -invariant Lie algebra valued 1-cocycles, yielding a reasonably explicit formula for a vector in the tangent space at $[\rho]$. Using this formula, it is then possible to compute an expression for the Goldman symplectic form in shear-shape coordinates. It remains to understand precisely how \mathcal{O} interacts with the various natural symplectic forms on \mathcal{PM}_g , \mathcal{QM}_g , and the (degenerate) symplectic forms on strata, a question that is made technical by the lack of regularity of \mathcal{O} .

Measures. To each $\text{PSL}_2 \mathbb{R}$ -invariant ergodic probability measure on \mathcal{QM}_g , pushforward along \mathcal{O}^{-1} produces a P -invariant ergodic probability measure on \mathcal{PT}_g (and vice-versa). An important class of such measures on the singular flat side is furnished by the Masur-Veech measure $\mu_{\underline{\kappa}}$ on a component of a stratum $\mathcal{Q}^1 \mathcal{M}_g(\underline{\kappa})$. In [CFa], the authors give a geometric description of $\nu_{\underline{\kappa}} := \mathcal{O}_*^{-1}(\mu_{\underline{\kappa}})$ on the corresponding “stratum” $\mathcal{P}^1 \mathcal{M}_g^{\text{reg}}(\underline{\kappa})$, which parallels [Mir08, Theorem 1.4] that on the principal stratum, $\nu_{\underline{\kappa}}$ disintegrates into the Weil-Petersson measure on hyperbolic surfaces and Thurston measure on laminations (up to a normalization factor). We give an outline of the various ingredients required to make the analogous statement for $\nu_{\underline{\kappa}}$ with $\underline{\kappa}$ arbitrary.

As discussed in Section 9.3, the piecewise-integral-linear (PIL) structure on $\mathcal{SH}(\lambda)$ endows it with an integer lattice and distinguished measure in the class of Lebesgue. Indeed, for each filling λ , the integer

lattice in $\mathcal{SH}^+(\lambda)$ restricts to an integer lattice on the fiber $\mathcal{H}^+(\lambda)$ over the empty arc system due to integrality of the equations defining the piecewise-linear structure of $\mathcal{B}(S \setminus \lambda)$. The empty arc system corresponds to the set of X on which λ is regular, and so the PIL structure induces a measure (in the class of Lebesgue) on this regular locus.

We identify the kernel of the Goldman symplectic form restricted to regular loci as tangent to certain “hyperbolic Schiffer deformations” associated to each even-gon in the complement of λ . These directions admit explicit descriptions as weight systems on a snug train track for λ [BW17, Appendix] which can be reinterpreted as 1-forms on regular loci obtained as the differentials of coordinate functions. Using our formula for the Goldman symplectic form restricted to regular loci, we identify the pullback of the Lebesgue measure on the fiber $\mathcal{H}(\lambda)$ over the empty arc system with an analytic volume form obtained as a wedge power of the restricted symplectic form then wedged together with the distinguished 1-forms associated to the kernel.

Using snug train tracks, one can define a $\underline{\kappa}$ -Thurston measure on the space $\mathcal{ML}(\underline{\kappa})$ of polygonal measured laminations of a given topological type. While this is essentially Lebesgue measure in train track coordinates for the “measure + Hausdorff” topology, it is not locally finite in the usual topology on $\mathcal{ML}(S)$. We then construct natural train track coordinate charts that give local measurable trivializations of $\mathcal{P}^1\mathcal{M}_g(\underline{\kappa})$ and exhibit $\nu_{\underline{\kappa}}$ as the product of $\underline{\kappa}$ -Thurston measure and the Weil-Petersson type volume form.

Counting. The integral points of the PIL structure on $\mathcal{SH}^+(\lambda)$ correspond to integer multicurves transverse to λ , so when λ is itself a multicurve, integral points correspond to square-tiled surfaces. Using our coordinates for $\mathcal{F}^{uu}(\lambda)$, one can recover the leading coefficient for the polynomial counting the number of square-tiled surfaces with given horizontal curve of bounded area (which was originally computed in [AH20a, DGZZ20]). In particular, since renormalized lattice point counts equidistribute to Lebesgue measure, the coefficient in question can be identified as the Lebesgue measure of (a torus bundle over) the portion of the combinatorial moduli space $\mathcal{B}(S \setminus \lambda) / \text{Mod}(S \setminus \lambda)$ with controlled boundary lengths. Using the convergence of the renormalized Weil-Petersson form to the Kontsevich form [Mon09b] (which on maximal cells induces the Lebesgue volume), this framework can also be manipulated to relate counts of curves on hyperbolic surfaces and intersection numbers on $\mathcal{M}_{g,n}$.

Moreover, since our coordinates also record the singularity type of the associated differential, the authors expect that the same strategy can also be used to count square-tiled surfaces in a given stratum with fixed horizontal multicurve. The leading coefficient should then be the Lebesgue volume of (a torus bundle over) a compact part of the appropriate subcomplex of the combinatorial moduli space. The relationships with hyperbolic geometry and intersection theory in these settings are more subtle, however, due in part to the non-transversality of the metric residue condition cutting out $\mathcal{B}(S \setminus \lambda)$ and the combinatorial conditions specifying strata. The authors hope to investigate these properties more fully in future work.

Expanding horospheres. Counting problems for square-tiled surfaces/curves on hyperbolic surfaces are intricately related to the equidistribution of L -level sets for the intersection number with/hyperbolic length of laminations as one takes $L \rightarrow \infty$. When λ is a multicurve, the equidistribution of such “expanding horospheres” to the Masur–Veech measure on the principal stratum of $\mathcal{Q}^1\mathcal{M}_g$ / the pullback by \mathcal{O} of this measure on $\mathcal{P}^1\mathcal{M}_g$ (sometimes called *Mirzakhani measure*) was established in [Mir07, AH20b, Liu20] using the geometry of the (symmetrized) Lipschitz metric, the non-divergence of the earthquake flow, and a no-escape-of-mass argument. On the other end of the spectrum, the equidistribution of expanding horospheres for maximal λ to $\mathcal{Q}^1\mathcal{M}_g$ can be proven using a standard “thickening plus mixing” argument from homogeneous dynamics; in the flat setting this is implicit in the work of Lindenstrauss and Mirzakhani [LM08], and was recently generalized in [For20, Theorem 1.6] using different methods. Equidistribution in the hyperbolic setting then follows from the continuity results described above.

Using our extension of Mirzakhani’s conjugacy (and the continuity results described above), the same “thickening plus mixing” technique can be used to prove that expanding horospheres based at any λ equidistribute to the Mirzakhani measure on $\mathcal{P}^1\mathcal{M}_g$. Moreover, an analogous result should also hold for strata: intersections of expanding horospheres based at λ and the regular locus should equidistribute to the pullback to $\mathcal{P}^1\mathcal{M}_g$ of the Masur–Veech measure for a component of $\mathcal{Q}^1\mathcal{M}_g(\underline{\kappa})$.

REFERENCES

- [AD20] D. Alessandrini and V. Disarlo, *Generalized stretch lines for surfaces with boundary*, Preprint, [arXiv:1911.10431](https://arxiv.org/abs/1911.10431), 2020.
- [AH20a] F. Arana-Herrera, *Counting square-tiled surfaces with prescribed real and imaginary foliations and connections to Mirzakhani's asymptotics for simple closed hyperbolic geodesics*, J. Mod. Dyn. **16** (2020), 81–107.
- [AH20b] ———, *Equidistribution of families of expanding horospheres on moduli spaces of hyperbolic surfaces*, Geom. Dedicata (2020).
- [AHW22] F. Arana-Herrera and A. Wright, *The asymmetry of Thurston's earthquake flow*, Preprint, [2201.04077](https://arxiv.org/abs/2201.04077), 2022.
- [BD17] F. Bonahon and G. Dreyer, *Hitchin characters and geodesic laminations*, Acta Math. **218** (2017), no. 2, 201–295.
- [BE88] B. Bowditch and D. B. A. Epstein, *Natural triangulations associated to a surface*, Topology **27** (1988), no. 1, 91–117.
- [Bon96] F. Bonahon, *Shearing hyperbolic surfaces, bending pleated surfaces and Thurston's symplectic form*, Ann. Fac. Sci. Toulouse Math. (6) **5** (1996), no. 2, 233–297.
- [Bon97a] ———, *Geodesic laminations with transverse Hölder distributions*, Ann. Sci. École Norm. Sup. (4) **30** (1997), no. 2, 205–240.
- [Bon97b] ———, *Transverse Hölder distributions for geodesic laminations*, Topology **36** (1997), no. 1, 103–122.
- [BS85] J. S. Birman and C. Series, *Geodesics with bounded intersection number on surfaces are sparsely distributed*, Topology **24** (1985), no. 2, 217–225.
- [Bus10] P. Buser, *Geometry and spectra of compact Riemann surfaces*, Modern Birkhäuser Classics, Birkhäuser Boston, Ltd., Boston, MA, 2010, Reprint of the 1992 edition.
- [BW17] F. Bonahon and H. Wong, *Representations of the Kauffman bracket skein algebra II: Punctured surfaces*, Algebr. Geom. Topol. **17** (2017), no. 6, 3399–3434.
- [CB88] A. J. Casson and S. A. Bleiler, *Automorphisms of surfaces after Nielsen and Thurston*, London Mathematical Society Student Texts, vol. 9, Cambridge University Press, Cambridge, 1988.
- [CEG06] R. D. Canary, D. B. A. Epstein, and P. L. Green, *Notes on notes of Thurston*, Fundamentals of hyperbolic geometry: selected expositions, London Math. Soc. Lecture Note Ser., vol. 328, Cambridge Univ. Press, Cambridge, 2006, With a new foreword by Canary, pp. 1–115.
- [CFa] A. Calderon and J. Farre, *Continuity of the orthogeodesic foliation and the geometry of train tracks*, In preparation.
- [CFb] ———, *Equidistribution on moduli spaces of flat and hyperbolic surfaces, with application to Mirzakhani's twist torus conjecture*, In preparation.
- [CM14] D. Chen and M. Möller, *Quadratic differentials in low genus: exceptional and non-varying strata*, Ann. Sci. Éc. Norm. Supér.(4) **47** (2014), no. 2, 309–369.
- [CSW20] J. Chaika, J. Smillie, and B. Weiss, *Tremors and horocycle dynamics on the moduli space of translation surfaces*, Preprint, [arXiv:2004.04027](https://arxiv.org/abs/2004.04027), 2020.
- [DGZZ20] V. Delecroix, E. Goujard, P. Zograf, and A. Zorich, *Masur-Veech volumes, frequencies of simple closed geodesics and intersection numbers of moduli spaces of curves*, Preprint, [arXiv:2011.05306](https://arxiv.org/abs/2011.05306), 2020.
- [Do08] N. Do, *Intersection theory on moduli spaces of curves via hyperbolic geometry*, Ph.D. thesis, The University of Melbourne, 2008.
- [Dum09] D. Dumas, *Complex projective structures*, Handbook of Teichmüller theory. Vol. II, IRMA Lect. Math. Theor. Phys., vol. 13, Eur. Math. Soc., Zürich, 2009, pp. 455–508. MR 2497780
- [Dum15] ———, *Skinning maps are finite-to-one*, Acta Math. **215** (2015), no. 1, 55–126. MR 3413977
- [EM06] D. B. A. Epstein and A. Marden, *Convex hulls in hyperbolic space, a theorem of Sullivan, and measured pleated surfaces*, Fundamentals of hyperbolic geometry: selected expositions, London Math. Soc. Lecture Note Ser., vol. 328, Cambridge Univ. Press, Cambridge, 2006, pp. 117–266.
- [EM18] A. Eskin and M. Mirzakhani, *Invariant and stationary measures for the $SL(2, \mathbb{R})$ action on moduli space*, Publ. Math. Inst. Hautes Études Sci. **127** (2018), 95–324.
- [EMM15] A. Eskin, M. Mirzakhani, and A. Mohammadi, *Isolation, equidistribution, and orbit closures for the $SL(2, \mathbb{R})$ action on moduli space*, Ann. of Math. (2) **182** (2015), no. 2, 673–721.
- [FK02] R. Feres and A. Katok, *Ergodic theory and dynamics of G -spaces (with special emphasis on rigidity phenomena)*, Handbook of dynamical systems, Vol. 1A, North-Holland, Amsterdam, 2002, pp. 665–763.
- [For20] G. Forni, *Limits of geodesic push-forwards of horocycle invariant measures*, Preprint, [arXiv:1908.11037](https://arxiv.org/abs/1908.11037), 2020.
- [Gel14] T. Gelander, *Lectures on lattices and locally symmetric spaces*, Preprint, [arXiv:1402.0962](https://arxiv.org/abs/1402.0962), 2014.
- [GM91] F. P. Gardiner and H. Masur, *Extremal length geometry of Teichmüller space*, Complex Variables Theory Appl. **16** (1991), no. 2–3, 209–237.
- [Gup17] S. Gupta, *Harmonic maps and wild Teichmüller spaces*, Preprint, [arXiv:2011.05306](https://arxiv.org/abs/2011.05306), 2017.
- [HM79] J. Hubbard and H. Masur, *Quadratic differentials and foliations*, Acta Math. **142** (1979), no. 3–4, 221–274.
- [HT] C. Horbez and J. Tao, *Nielsen-Thurston classification and isometries of the Thurston metric*, In preparation.
- [HTTW95] Z. Han, L. Tam, A. Treibergs, and T. Wan, *Harmonic maps from the complex plane into surfaces with nonpositive curvature*, Comm. Anal. Geom. **3** (1995), no. 1–2, 85–114.
- [Ker83] S. P. Kerckhoff, *The Nielsen realization problem*, Ann. of Math. (2) **117** (1983), no. 2, 235–265.
- [KZ03] M. Kontsevich and A. Zorich, *Connected components of the moduli spaces of Abelian differentials with prescribed singularities*, Invent. Math. **153** (2003), no. 3, 631–678.
- [Lan08] E. Lanneau, *Connected components of the strata of the moduli spaces of quadratic differentials*, Ann. Sci. l'Éc. Norm. Sup. **41** (2008), no. 1, 1–56.

- [Lev83] G. Levitt, *Foliations and laminations on hyperbolic surfaces*, Topology **22** (1983), no. 2, 119–135.
- [Liu20] M. Liu, *Length statistics of random multicurves on closed hyperbolic surfaces*, Preprint, [arXiv:1912.11155](https://arxiv.org/abs/1912.11155), 2020.
- [LM08] E. Lindenstrauss and M. Mirzakhani, *Ergodic theory of the space of measured laminations*, Int. Math. Res. Not. IMRN (2008), no. 4, Art. ID rnm126, 49.
- [Luo07] F. Luo, *On Teichmüller spaces of surfaces with boundary*, Duke Math. J. **139** (2007), no. 3, 463–482.
- [Mas82] H. Masur, *Interval exchange transformations and measured foliations*, Ann. Math. **115** (1982), no. 1, 169–200.
- [Min92] Y. N. Minsky, *Harmonic maps, length, and energy in Teichmüller space*, J. Differ. Geom. **35** (1992), 151–217.
- [Mir07] M. Mirzakhani, *Random hyperbolic surfaces and measured laminations*, In the tradition of Ahlfors-Bers. IV, Contemp. Math., vol. 432, Amer. Math. Soc., Providence, RI, 2007, pp. 179–198.
- [Mir08] _____, *Ergodic theory of the earthquake flow*, Int. Math. Res. Not. IMRN (2008), no. 3, Art. ID rnm116, 39.
- [Mon09a] G. Mondello, *Riemann surfaces, ribbon graphs and combinatorial classes*, Handbook of Teichmüller theory. Vol. II, IRMA Lect. Math. Theor. Phys., vol. 13, Eur. Math. Soc., Zürich, 2009, pp. 151–215.
- [Mon09b] _____, *Triangulated Riemann surfaces with boundary and the Weil-Petersson Poisson structure*, J. Differential Geom. **81** (2009), no. 2, 391–436.
- [MW02] Y. N. Minsky and B. Weiss, *Nondivergence of horocyclic flows on moduli space*, J. Reine Angew. Math. **552** (2002), 131–177.
- [MW14] _____, *Cohomology classes represented by measured foliations, and Mahler’s question for interval exchanges*, Ann. Sci. Éc. Norm. Supér. (4) **47** (2014), no. 2, 245–284.
- [Pap86] A. Papadopoulos, *Geometric intersection functions and Hamiltonian flows on the space of measured foliations on a surface*, Pacific J. Math. **124** (1986), no. 2, 375–402.
- [Pap91] _____, *On Thurston’s boundary of Teichmüller space and the extension of earthquakes*, Topology Appl. **41** (1991), no. 3, 147–177.
- [PH92] R. C. Penner and J. L. Harer, *Combinatorics of train tracks*, Annals of Mathematics Studies, vol. 125, Princeton University Press, Princeton, NJ, 1992.
- [PT07] A. Papadopoulos and G. Théret, *On Teichmüller’s metric and Thurston’s asymmetric metric on Teichmüller space*, Handbook of Teichmüller theory. Vol. I, IRMA Lect. Math. Theor. Phys., vol. 11, Eur. Math. Soc., Zürich, 2007, pp. 111–204.
- [PW22] H. Pan and M. Wolf, *Ray structures on Teichmüller space*, Preprint, [arXiv:2206.01371](https://arxiv.org/abs/2206.01371), 2022.
- [PY17] A. Papadopoulos and S. Yamada, *Deforming hyperbolic hexagons with applications to the arc and the Thurston metrics on Teichmüller spaces*, Monatsh. Math. **182** (2017), no. 4, 913–939.
- [SB01] Y. Sözen and F. Bonahon, *The Weil-Petersson and Thurston symplectic forms*, Duke Math. J. **108** (2001), no. 3, 581–597.
- [Sul85] D. Sullivan, *Quasiconformal homeomorphisms and dynamics. II. Structural stability implies hyperbolicity for Kleinian groups*, Acta Math. **155** (1985), no. 3–4, 243–260.
- [Thu82] W. P. Thurston, *The geometry and topology of 3-manifolds*, 1982, Princeton University Lecture Notes.
- [Thu97] _____, *Three-dimensional geometry and topology. Vol. 1*, Princeton Mathematical Series, vol. 35, Princeton University Press, Princeton, NJ, 1997, Edited by Silvio Levy.
- [Thu98] _____, *Minimal stretch maps between hyperbolic surfaces*, Preprint, [arXiv:9801039](https://arxiv.org/abs/9801039), 1998.
- [Ush99] A. Ushijima, *A canonical cellular decomposition of the Teichmüller space of compact surfaces with boundary*, Comm. Math. Phys. **201** (1999), no. 2, 305–326.
- [Vee82] W. Veech, *Gauss measures for transformations on the space of interval exchange maps*, Ann. Math. **115** (1982), 201–242.
- [Wol83] S. Wolpert, *On the symplectic geometry of deformations of a hyperbolic surface*, Ann. of Math. (2) **117** (1983), no. 2, 207–234.
- [Wri18] A. Wright, *Mirzakhani’s work on earthquake flow*, Preprint, [arXiv:1810.07571](https://arxiv.org/abs/1810.07571), 2018.
- [Wri20] _____, *A tour through Mirzakhani’s work on moduli spaces of Riemann surfaces*, Bull. Amer. Math. Soc. (N.S.) **57** (2020), no. 3, 359–408.
- [ZB04] X. Zhu and F. Bonahon, *The metric space of geodesic laminations on a surface. I*, Geom. Topol. **8** (2004), 539–564.

INDEX

α	60	$H(v, w)$	72
$\underline{\alpha}$	16	Hexagons	52
$\alpha(q)$	41	Hexagonal hulls	72
$A(q)$	43		
$A(X)$	52	$I_\lambda(q)$	43
$ \mathcal{A}_{\text{fill}}(\Sigma, \partial\Sigma) _{\mathbb{R}}$	19	$\text{inj}_\lambda(X)$	61
$A(\vec{s})$	64	k_b^d	61
$A(\vec{\alpha}_1, u)$	70	$\ell_X(b)$	17
Admissible routes	74	$\ell_s(\alpha_1)$	69
$\mathcal{B}(S \setminus \lambda)$	28	$\hat{\lambda} \cup \hat{\underline{\alpha}}$	25
Binding pairs of laminations	5	$\Lambda_{v,w}$	53
Boundary leaves	41	Metric residue	13
c_α	43, 52	$\mathcal{MF}(\lambda)$	3
Crowned hyperbolic surface	12		
δ_\pm	72	$\mathcal{N}_\epsilon(\lambda)$	16
$D_\lambda(X)$	62	$n_0(\lambda)$	25
$d_A(\vec{\alpha}_1, u)$	69	\mathcal{O}	5
$\Delta(\lambda)$	34	$\mathcal{O}_{\partial Y}(Y)$	15
$ \Delta(\lambda) _\pm$	48	\mathcal{O}_λ	16
$\partial_\lambda H_v$	52	Orientations of $\lambda \cup \underline{\alpha}$	25
Deflation	18	Orientation double cover	25
Depth	61		
Dilation rays	81	\mathcal{PT}_g and $\mathcal{P}^1\mathcal{T}_g$	5
Dual arcs	16	$\mathcal{P}^1\mathcal{M}_g^{\text{reg}}(\kappa)$	1
$\varphi(\vec{s})$	64	Piecewise-integral-linear	39
φ_H	64	Pointed geodesic	52
$\varphi_{v,w}^r$	67		
φ_{p_v, p_w}	73	\mathcal{QT}_g and $\mathcal{Q}^1\mathcal{T}_g$	5
φ_s	76	$\mathcal{Q}^1\mathcal{M}_g^{\text{nsc}}(\kappa)$	1
Φ_{p_v}	78		
$\mathcal{F}^{uu}(\lambda)$	3	$r_b(d)$	61
F_λ	3	$\text{res}(\mathcal{C})$	13
$f_{X,s}(\vec{s})$	63	$\text{res}_A(\mathcal{C})$	19
$f_{X,s}(\vec{\alpha}_1, u)$	69	ρ_s	78
(g, p_v)	52		
g_v^w, g_w^v	53	\vec{s}	64
$\mathcal{G}(\tilde{X})$	78	\mathfrak{s}	60
Geometric standard smoothings	37	$\ \mathfrak{s}\ _{\vec{s}}$	63
Geometric train tracks	16	$\sigma_\lambda(X)$	56
\mathcal{H}	52	$\Sigma_{g,b}^{\{\underline{\alpha}\}}$	12
$h_A(s)$	63	$\widehat{\Sigma}_{g,b}^{\{\underline{\alpha}\}}$	13
$\mathcal{H}(\lambda)$	25	Singular leaves	41
$\mathcal{H}^+(\lambda)$	34	Shape-shifting cocycles	74
(H0), (H1), (H2)	25	Shear-shape cocycles	26, 29
$H^1(\hat{N}, \partial\hat{N}; \mathbb{R})^-$	25	$\mathcal{SH}(\lambda)$	26
$H^1(\hat{N}_{\underline{\alpha}}, \partial\hat{N}_{\underline{\alpha}}; \mathbb{R})^-$	26	$\mathcal{SH}^+(\lambda)$	35
		$\mathcal{SH}^\circ(\lambda; \underline{\alpha})$	26

$\mathcal{SH}(\lambda; A)$	26	Ties	16
(SH0), (SH1), (SH2), (SH3)	29	Transverse cocycles	25
Simple pairs	53	Transverse pairs	57
Snug neighborhoods	25	Tremors, $\text{trem}_\mu(q)$	48
Spine	15	Valence, $\text{val}(x)$	15
$\text{Sp}(Y)$	15	$\omega_{\mathcal{H}}$	34
$\text{Sp}_k(Y)$	15	$\omega_{\mathcal{SH}}$	35
Sp^0	15	$w_{\underline{\alpha}}(\sigma)$	38
Standard transversals	25	$\Xi(q)$	41
Standard smoothings	37	$\chi(\lambda)$	14
Stretchquakes	82	$X_{\mathfrak{s}}$	78
$\tau_{\underline{\alpha}}$	37	$[z]_+$	44
τ_{\max}	61		
$T^* \setminus H^*$	42		
Thick with respect to v and w	74		

Aaron Calderon, DEPARTMENT OF MATHEMATICS, UNIVERSITY OF CHICAGO

E-mail address: aaroncalderon@uchicago.edu

James Farre, DEPARTMENT OF MATHEMATICS, UNIVERSITÄT HEIDELBERG

E-mail address: jfarre@mathi.uni-heidelberg.de